
Quadratic optimal functional quantization of stochastic processes and numerical applications

Gilles Pagès¹

Laboratoire de Probabilités et Modèles aléatoires, UMR 7599, Université Paris 6, case 188, 4, pl. Jussieu, F-75252 Paris Cedex 5, France. gpa@ccr.jussieu.fr

Summary. In this paper, we present an overview of the recent developments of functional quantization of stochastic processes, with an emphasis on the quadratic case. Functional quantization is a way to approximate a process, viewed as a Hilbert-valued random variable, using a nearest neighbour projection on a finite codebook. A special emphasis is made on the computational aspects and the numerical applications, in particular the pricing of some path-dependent European options.

1 Introduction

Functional quantization is a way to discretize the path space of a stochastic process. It has been extensively investigated since the early 2000's by several authors (see among others [29], [31], [12], [9], [30], etc). It first appeared as a natural extension of the Optimal Vector Quantization theory of (finite-dimensional) random vectors which finds its origin in the early 1950's for signal processing (see [15] or [17]).

Let us consider a Hilbertian setting. One considers a random vector X defined on a probability space $(\Omega, \mathcal{A}, \mathbb{P})$ taking its values in a separable Hilbert space $(H, (\cdot|\cdot)_H)$ (equipped with its natural Borel σ -algebra) and satisfying $\mathbb{E}|X|^2 < +\infty$. When H is an Euclidean space (\mathbb{R}^d) , one speaks about *Vector Quantization*. When H is an infinite dimensional space like $L^2_T := L^2([0, T], dt)$ (endowed with the usual Hilbertian norm $|f|_{L^2_T} := (\int_0^T f^2(t)dt)^{\frac{1}{2}}$) one speaks of *functional quantization* (denoted L^2_T from now on). A (bi-measurable) stochastic process $(X_t)_{t \in [0, T]}$ defined on $(\Omega, \mathcal{A}, \mathbb{P})$ satisfying $|X(\omega)|_{L^2_T} < +\infty$ $\mathbb{P}(d\omega)$ -a.s. can always be seen, once possibly modified on a \mathbb{P} -negligible set, as an L^2_T -valued random variable. Although we will focus on the Hilbertian framework, other choices are possible for H , in particular some more general Banach settings like $L^p([0, T], dt)$ or $\mathcal{C}([0, T], \mathbb{R})$ spaces.

This paper is organized as follows: in Sections 2 we introduce quadratic quantization in a Hilbertian setting. In Section 3, we focus on optimal quantization, including some extensions to non quadratic quantization. Section 4 is

devoted to some quantized cubature formulae. Section 5 provides some classical background on the quantization rate in finite dimension. Section 7 deals with functional quantizations of Gaussian processes, like the Brownian motion, with a special emphasis on the numerical aspects. We present here what is, to our guess, the first large scale numerical optimization of the quadratic quantization of the Brownian motion. We compare it to the optimal product quantization, formerly investigated in [44]. In section, we propose a constructive approach to the functional quantization of scalar or multidimensional diffusions (in the Stratanovich sense). In Section 9, we show how to use functional quantization to price path-dependent options like Asian options (in a heston stochastic volatility model). We conclude by some recent results showing how to derive universal (often optimal) functional quantization rate from time regularity of a process in Section 10 and by a few clues in Section 11 about the specific methods that produce some lower bounds (this important subject as many others like the connections with small deviation theory is not treated in this numerically oriented overview. As concerns statistical applications of functional quantization we refer to [53, 54].

NOTATIONS. • $a_n \approx b_n$ means $a_n = O(b_n)$ and $b_n = O(a_n)$; $a_n \sim b_n$ means $a_n = b_n + o(a_n)$.

- If $X : (\Omega, \mathcal{A}, \mathbb{P}) \rightarrow (H, |\cdot|_H)$ (Hilbert space), then $\|X\|_2 = (\mathbb{E}|X|_H^2)^{\frac{1}{2}}$.
- $\lfloor x \rfloor$ denotes the integral part of the real x .

2 What is quadratic functional quantization?

Let $(H, (\cdot|\cdot)_H)$ denote a separable Hilbert space. Let $X \in L^2_H(\mathbb{P})$ i.e. a random vector $X : (\Omega, \mathcal{A}, \mathbb{P}) \mapsto H$ (H is endowed with its Borel σ -algebra) such that $\mathbb{E}|X|_H^2 < +\infty$. An N -quantizer (or N -codebook) is defined as a subset

$$\Gamma := \{x_1, \dots, x_N\} \subset H$$

with $\text{card}\Gamma = N$. In numerical applications, Γ is also called *grid*. Then, one can *quantize* (or simply discretize) X by $q(X)$ where $q : H \mapsto \Gamma$ is a Borel function. It is straightforward that

$$\forall \omega \in \Omega, \quad |X(\omega) - q(X(\omega))|_H \geq d(X(\omega), \Gamma) = \min_{1 \leq i \leq N} |X(\omega) - x_i|_H$$

so that the best pointwise approximation of X is provided by considering for q a nearest neighbour projection on Γ , denoted Proj_Γ . Such a projection is in one-to-one correspondence with the Voronoi partitions (or diagrams) of H induced by Γ i.e. the Borel partitions of H satisfying

$$C_i(\Gamma) \subset \left\{ \xi \in H : |\xi - x_i|_H = \min_{1 \leq j \leq N} |\xi - x_j|_H \right\} = \overline{C}_i(\Gamma), \quad i = 1, \dots, N,$$

where $\overline{C}_i(\Gamma)$ denotes the closure of $C_i(\Gamma)$ in H (this heavily uses the Hilbert structure). Then

$$\text{Proj}_\Gamma(\xi) := \sum_{i=1}^N x_i \mathbf{1}_{C_i(\Gamma)}(\xi)$$

is a nearest neighbour projection on Γ . These projections only differ on the boundaries of the *Voronoi cells* $C_i(\Gamma)$, $i = 1, \dots, N$. All Voronoi partitions have the same boundary contained in the union of the median hyperplanes defined by the pairs (x_i, x_j) , $i \neq j$. Figure 2 represents the Voronoi diagram defined by a (random) 10-tuple in \mathbb{R}^2 . Then, one defines a *Voronoi*




Fig. 1. A 2-dimensional 10-quantizer $\Gamma = \{x_1, \dots, x_{10}\}$ and its Voronoi diagram.

N -quantization of X by setting for every $\omega \in \Omega$,

$$\widehat{X}^\Gamma(\omega) := \text{Proj}_\Gamma(X(\omega)) = \sum_{i=1}^N x_i \mathbf{1}_{C_i(\Gamma)}(X(\omega)).$$

One clearly has, still for every $\omega \in \Omega$, that

$$|X(\omega) - \widehat{X}^\Gamma(\omega)|_H = \text{dist}_H(X(\omega), \Gamma) = \min_{1 \leq i \leq N} |X(\omega) - x_i|_H.$$

The mean (quadratic) quantization error is then defined by

$$e(\Gamma, X, H) = \|X - \widehat{X}^\Gamma\|_2 = \sqrt{\mathbb{E} \left(\min_{1 \leq i \leq N} |X - x_i|_H^2 \right)}. \quad (1)$$

The distribution of \widehat{X}^Γ as a random vector is given by the N -tuple $(\mathbb{P}(X \in C_i(\Gamma)))_{1 \leq i \leq N}$ of the Voronoi cells. This distribution clearly depends on the choice of the Voronoi partition as emphasized by the following elementary situation: if $H = \mathbb{R}$, the distribution of X is given by $\mathbb{P}_X = \frac{1}{3}(\delta_0 + \delta_{1/2} + \delta_1)$,

$N = 2$ and $\Gamma = \{0, 1\}$ since $1/2 \in \partial C_0(\Gamma) \cap \partial C_1(\Gamma)$. However, if \mathbb{P}_x weights no hyperplane, the distribution of \hat{X}^Γ depends only on Γ .

As concerns terminology, *Vector Quantization* is concerned with the finite dimensional case – when $\dim H < +\infty$ – and is a rather old story, going back to the early 1950's when it was designed in the field of signal processing and then mainly developed in the community of Information Theory. The term *functional quantization*, probably introduced in [41, 29], deals with the infinite dimensional case including the more general Banach-valued setting. The term “functional” comes from the fact that a typical infinite dimensional Hilbert space is the function space $H = L^2_T$. Then, any (bi-measurable) process $X : ([0, T] \times \Omega, \text{Bor}([0, T]) \otimes \mathcal{A}) \rightarrow (\mathbb{R}, \text{Bor}(\mathbb{R}))$ can be seen as a random vector taking values in the set of Borel functions on $[0, T]$. Furthermore, $((t, \omega) \mapsto X_t(\omega)) \in L^2(dt \otimes d\mathbb{P})$ if and only if $(\omega \mapsto X(\omega)) \in L^2_H(\mathbb{P})$ since

$$\int_{[0, T] \times \Omega} X_t^2(\omega) dt \mathbb{P}(d\omega) = \int_{\Omega} \mathbb{P}(d\omega) \int_0^T X_t^2(\omega) dt = \mathbb{E} \|X\|_{L^2_T}^2.$$

3 Optimal (quadratic) quantization

At this stage we are lead to wonder whether it is possible to design some optimally fitted grids to a given distribution \mathbb{P}_x *i.e.* which induce the lowest possible mean quantization error among all grids of size at most N . This amounts to the following optimization problem

$$e_N(X, H) := \inf_{\Gamma \subset H, \text{card}(\Gamma) \leq N} e(\Gamma, X, H). \quad (2)$$

It is convenient at this stage to make a correspondence between quantizers of size at most N and N -tuples of H^N : to any N -tuple $x := (x_1, \dots, x_N)$ corresponds a quantizer $\Gamma := \Gamma(x) = \{x_i, i = 1, \dots, N\}$ (of size at most N). One introduces the quadratic distortion, denoted D_N^X , defined on H^N as a (symmetric) function by

$$D_N^X : H^N \longrightarrow \mathbb{R}_+ \\ (x_1, \dots, x_N) \longmapsto \mathbb{E} \left(\min_{1 \leq i \leq N} |X - x_i|_H^2 \right).$$

Note that, combining (1) and the definition of the distortion, shows that

$$D_N^X(x_1, \dots, x_N) = \mathbb{E} \left(\min_{1 \leq i \leq N} |X - x_i|_H^2 \right) = \mathbb{E} (d(X, \Gamma(x))^2) = \|X - \hat{X}^{\Gamma(x)}\|_2^2$$

so that,




Fig. 2. Two N -quantizers (and their Voronoi diagram) related to bi-variate normal distribution $\mathcal{N}(0; I_2)$ ($N = 500$); which one is the best?

$$e_N(X, H) = \inf_{(x_1, \dots, x_N) \in H^N} \sqrt{D_N^X(x_1, \dots, x_N)}.$$

The following proposition shows the existence of an optimal N -tuple $x^{(N,*)} \in H^N$ such that $e_N(X, H) = \sqrt{D_N^X(x^{(N,*)})}$. The corresponding optimal quantizer at level N is denoted $\Gamma^{(N,*)} := \Gamma(x^{(N,*)})$. In finite dimension we refer to [49] (1982) and in infinite dimension to [7] (1988) and [48] (1990); one may also see [39], [17] and [29]. For recent developments on existence and pathwise regularity of optimal quantizer see [20].

Proposition 1. (a) The function D_N^X is lower semi-continuous for the product weak topology on H^N .

(b) The function D_N^X reaches a minimum at a N -tuple $x^{(N,*)}$ (so that $\Gamma^{(N,*)}$ is an optimal quantizer at level N).

– If $\text{card}(\text{supp}(\mathbb{P}_X)) \geq N$, the quantizer has full size N (i.e. $\text{card}(\Gamma^{(N,*)}) = N$) and $e_N(X, H) < e_{N-1}(X, H)$.

– If $\text{card}(\text{supp}(\mathbb{P}_X)) \leq N$, $e_N(X, H) = 0$.

Furthermore $\lim_N e_N(X, H) = 0$.

(c) Any optimal (Voronoi) quantization at level N , $\hat{X}^{\Gamma^{(N,*)}}$ satisfies

$$\hat{X}^{\Gamma^{(N,*)}} = \mathbb{E}(X \mid \sigma(\hat{X}^{\Gamma^{(N,*)}})) \quad (3)$$

where $\sigma(\hat{X}^{\Gamma^{(N,*)}})$ denotes the σ -algebra generated by $\hat{X}^{\Gamma^{(N,*)}}$.

(d) Any optimal (quadratic) quantization at level N is a best least square (i.e. $L^2(\mathbb{P})$) approximation of X among all H -valued random variables taking at most N values:

$$e_N(X, H) = \|X - \hat{X}^{\Gamma^{(N,*)}}\|_2 = \min\{\|X - Y\|_2, Y : (\Omega, \mathcal{A}) \rightarrow H, \text{card}(Y(\Omega)) \leq N\}.$$

Proof (sketch of): (a) The claim follows from the l.s.c. of $\xi \mapsto |\xi|_H$ for the weak topology and Fatou's Lemma.

(b) One proceeds by induction on N . If $N = 1$, the optimal 1-quantizer is $x^{(N,*)} = \{\mathbb{E} X\}$ and $e_2(X, H) = \|X - \mathbb{E} X\|_2$.

Assume now that an optimal quantizer $x^{(N,*)} = (x_1^{(N,*)}, \dots, x_N^{(N,*)})$ does exist at level N .

– If $\text{card}(\text{supp}(\mathbb{P})) \leq N$, then the $N + 1$ -tuple $(x^{(N,*)}, x_N^{(N,*)})$ (among other possibilities) is also optimal at level $N + 1$ and $e_{N+1}(X, H) = e_N(X, H) = 0$.

– Otherwise, $\text{card}(\text{supp}(\mathbb{P})) \geq N + 1$, hence $x^{(N,*)}$ has pairwise distinct components and there exists $\xi_{N+1} \in \text{supp}(\mathbb{P}_X) \setminus \{x_i^{(N,*)}, i = 1, \dots, N\} \neq \emptyset$.

Then, with obvious notations,

$$D_{N+1}^X((x^{(N,*)}, \xi_{N+1})) < D_N^X(x^{(N,*)}).$$

Then, the set $F_{N+1} := \{x \in H^{N+1} \mid D_{N+1}^X(x) \leq D_{N+1}^X((x^{(N,*)}, \xi_{N+1}))\}$ is non empty, weakly closed since D_{N+1}^X is l.s.c.. Furthermore, it is bounded in H^{N+1} . Otherwise there would exist a sequence $x_{(m)} \in H^{N+1}$ such that $|x_{(m), i_m}|_H = \max_i |x_{(m), i}|_H \rightarrow +\infty$ as $m \rightarrow \infty$. Then, by Fatou's Lemma, one checks that

$$\liminf_{m \rightarrow \infty} D_{N+1}^X(x_{(m)}) \geq D_N^X(x^{(N,*)}) > D_{N+1}^X((x^{(N,*)}, \xi_{N+1})).$$

Consequently F_{N+1} is weakly compact and the minimum of D_{N+1}^X on F_{N+1} is clearly its minimum over the whole space H^{N+1} . In particular

$$e_{N+1}(X, H) \leq D_{N+1}^X((x^{(N,*)}, \xi_{N+1})) < e_N(X, H).$$

If $\text{card}(\text{supp}(\mathbb{P})) = N + 1$, set $x^{(N+1,*)} = \text{supp}(\mathbb{P})$ (as sets) so that $t X = \widehat{X}^{\Gamma^{(N+1,*)}}$ which implies $e_{N+1}(X, H) = 0$.

To establish that $e_N(X, H)$ goes to 0, one considers an everywhere dense sequence $(z_k)_{k \geq 1}$ in the separable space H . Then, $d(\{z_1, \dots, z_N\}, X(\omega))$ goes to 0 as $N \rightarrow \infty$ for every $\omega \in \Omega$. Furthermore, $d(\{z_1, \dots, z_N\}, X(\omega))^2 \leq |X(\omega) - z_1|_H^2 \in L^1(\mathbb{P})$. One concludes by the Lebesgue dominated convergence Theorem that $D_N^X(z_1, \dots, z_N)$ goes to 0 as $N \rightarrow \infty$.

(c) and (d) Temporarily set $\widehat{X}^* := \widehat{X}^{\Gamma^{(N,*)}}$ for convenience. Let $Y : (\Omega, \mathcal{A}) \rightarrow H$ be a random vector taking at most N values. Set $\Gamma := Y(\Omega)$. Since \widehat{X}^Γ is a Voronoi quantization of X induced by Γ ,

$$|X - \widehat{X}^\Gamma|_H = d(X, \Gamma) \leq |X - Y|_H$$

so that

$$\|X - \widehat{X}^\Gamma\|_2 \leq \|X - Y\|_2.$$

On the other hand, the optimality of $\Gamma^{(N,*)}$ implies

$$\|X - \widehat{X}^*\|_2 \leq \|X - \widehat{X}^\Gamma\|_2.$$

Consequently

$$\|X - \widehat{X}^*\|_2 \leq \min \{ \|X - Y\|_2, Y : (\Omega, \mathcal{A}) \rightarrow H, \text{card}(Y(\Omega)) \leq N \}.$$

The inequality holds as an equality since \widehat{X}^* takes at most N values. Furthermore, considering random vectors of the form $Y = g(\widehat{X})$ (which take at most as many values as the size of $\Gamma^{(N,*)}$) shows, going back to the very definition of conditional expectation, that $\widehat{X}^* = \mathbb{E}(X | \widehat{X}^*)$ \mathbb{P} -a.s. \diamond

Item (c) introduces a very important notion in (quadratic) quantization.

Definition 1. A quantizer $\Gamma \subset H$ is stationary (or self-consistent) if (there is a nearest neighbour projection such that $\widehat{X}^\Gamma = \text{Proj}_\Gamma(X)$ satisfying)

$$\widehat{X}^\Gamma = \mathbb{E}(X | \widehat{X}^\Gamma). \quad (4)$$

Note in particular that any stationary quantization satisfies $\mathbb{E}X = \mathbb{E}\widehat{X}^\Gamma$.

As shown by Proposition 1(c) any quadratic optimal quantizer at level N is stationary. Usually, at least when $d \geq 2$, there are other stationary quantizers: indeed, the distortion function D_N^X is $|\cdot|_H$ -differentiable at N -quantizers $x \in H^N$ with pairwise distinct components and

$$\nabla D_N^X(x) = 2 \left(\int_{C_i(x)} (x_i - \xi) \mathbb{P}_x(d\xi) \right)_{1 \leq i \leq N} = 2 \left(\mathbb{E}(\widehat{X}^{\Gamma(x)} - X) \mathbf{1}_{\{\widehat{X}^{\Gamma(x)} = x_i\}} \right)_{1 \leq i \leq N}.$$

hence, any critical points of D_N^X is a stationary quantizer.

Remarks and comments. • In fact (see Theorem 4.2, p. 38, [17]), the Voronoi partitions of $\Gamma^{(N,*)}$ always have a \mathbb{P}_X -negligible boundary so that (4) holds for *any* Voronoi diagram induced by Γ .

• The problem of the uniqueness of optimal quantizer (viewed as a set) is not mentioned in the above proposition. In higher dimension, this essentially never occurs. In one dimension, uniqueness of the optimal N -quantizer was first established in [14] with strictly log-concave density function. This was successively extended in [23] and [55] and lead to the following criterion (for more general “loss” functions than the square function):

If the distribution of X is *absolutely continuous with a log-concave density function*, then, for every $N \geq 1$, there exists only one stationary quantizer of size N , which turns out to be the optimal quantizer at level N .

More recently, a more geometric approach to uniqueness based on the Mountain Pass Lemma first developed in [26] and then generalized in [6]) provided a slight extension of the above criterion (in terms of loss functions).

This log-concavity assumption is satisfied by many families of probability distributions like the uniform distribution on compact intervals, the normal distributions, the gamma distributions. There are examples of distributions with a non log-concave density function having a unique optimal quantizer for every $N \geq 1$ (see *e.g.* the Pareto distribution in [16]). On the other hand simple examples of scalar distributions having multiple optimal quantizers at a given level can be found in [17].

- A stationary quantizer can be sub-optimal. This will be emphasized in Section 7 for the Brownian motion (but it is also true for finite dimensional Gaussian random vectors) where some families of sub-optimal quantizers – the *product quantizers* designed from the Karhunen-Loève basis – are stationary quantizers.

- For the uniform distribution over an interval $[a, b]$, there is a closed form for the optimal quantizer at level N given by $I^{(N,*)} = \{a + (2k - 1)\frac{b-a}{N}, k = 1, \dots, N\}$. This N -quantizer is optimal not only in the quadratic case but also for any L^r -quantization (see a definition further on). In general there is no such closed form, either in 1 or higher dimension. However, in [16] some semi-closed forms are obtained for several families of (scalar) distributions including the exponential and the Pareto distributions: all the optimal quantizers can be expressed using a single underlying sequence $(a_k)_{k \geq 1}$ defined by an induction $a_{k+1} = F(a_k)$.

- In one dimension, as soon as the optimal quantizer at level N is unique (as a set or as an N -tuple with increasing components), it is generally possible to compute it as the solution of the stationarity equation (3) either by a zero search (Newton-Raphson gradient descent) or a fixed point (like the specific Lloyd I procedure, see [24]) procedure.

- In higher dimension, deterministic optimization methods become intractable and one uses stochastic procedures to compute optimal quantizers. The main topic of this paper being functional quantization, we postponed the short overview on these aspects to Section 7, devoted to the optimal quantization of the Brownian motion. But it is to be noticed that all efficient optimization methods rely on the so-called *splitting method* which increases progressively the quantization level N . This method is directly inspired by the induction developed in the proof of claim (b) of Proposition 1 since one designs the starting value of the optimization procedure at size $N + 1$ by “merging” the optimized N -quantizer obtained at level N with one further point of \mathbb{R}^d , usually randomly sampled with respect to an appropriate distribution (see [43] for a discussion).

- As concerns functional quantization, *e.g.* $H = L^2_T$, there is a close connection between the regularity of optimal (or even stationary) quantizers and that of $t \mapsto X_t$ from $[0, T]$ into $L^2(\mathbb{P})$. Furthermore, as concerns optimal quantizers of

Gaussian processes, one shows (see [29]) that they belong to the reproducing space of their covariance operator, *e.g.* to the Cameron-Martin space $H^1 = \{\int_0^\cdot \dot{h}_s ds, \dot{h} \in L^2_T\}$ when $X = W$. Other properties of optimal quantization of Gaussian processes are established in [29].

Extensions to the $L^r(\mathbb{P})$ -quantization of random variables. In this paper, we focus on the purely quadratic framework (L^2_T and $L^2(\mathbb{P})$ -norms), essentially because it is a natural (and somewhat easier) framework for the computation of optimized grids for the Brownian motion and for some first applications (like the pricing of path-dependent options, see section 9). But a more general and natural framework is to consider the functional quantization of random vectors taking values in a separable Banach space $(E, |\cdot|_E)$. Let $X : (\Omega, \mathcal{A}, \mathbb{P}) \rightarrow (E, |\cdot|_E)$, such that $\mathbb{E}|X|_E^r < +\infty$ for some $r \geq 1$ (the case $0 < r < 1$ can also be taken in consideration).

The N -level $(L^r(\mathbb{P}), |\cdot|_E)$ -quantization problem for $X \in L^r_E(\mathbb{P})$ reads

$$e_{N,r}(X, E) := \inf \left\{ \|X - \hat{X}^\Gamma\|_r, \Gamma \subset E, \text{card}(\Gamma) \leq N \right\}.$$

The main examples for $(E, |\cdot|_E)$ are the non-Euclidean norms on \mathbb{R}^d , the functional spaces $L^p_T(\mu) := L^p([0, T], \mu(dt))$, $1 \leq p \leq \infty$, equipped with its usual norm, $(E, |\cdot|_E) = (\mathcal{C}([0, T]), \|\cdot\|_{\sup})$, etc. As concerns, the existence of an optimal quantizer, it holds true for reflexive Banach spaces (see Pärna (90)) and $E = L^1_T$, but otherwise it may fail even when $N = 1$ (see [20]). In finite dimension, the Euclidean feature is not crucial (see [17]). In the functional setting, many results originally obtained in a Hilbert setting have been extended to the Banach setting either for existence or regularity results (see [20]) or for rates see [10], [12], [30], [33].

4 Cubature formulae: conditional expectation and numerical integration

Let $F : H \rightarrow \mathbb{R}$ be a continuous functional (with respect to the norm $|\cdot|_H$) and let $\Gamma \subset H$ be an N -quantizer. It is natural to approximate $\mathbb{E}(F(X))$ by $\mathbb{E}(F(\hat{X}^\Gamma))$. This quantity $\mathbb{E}(F(\hat{X}^\Gamma))$ is simply the finite weighted sum

$$\mathbb{E}(F(\hat{X}^\Gamma)) = \sum_{i=1}^N F(x_i) \mathbb{P}(\hat{X}^\Gamma = x_i).$$

Numerical computation of $\mathbb{E}(F(\hat{X}^\Gamma))$ is possible as soon as $F(\xi)$ can be computed at any $\xi \in H$ and the distribution $(\mathbb{P}(\hat{X}^\Gamma = x_i))_{1 \leq i \leq N}$ of \hat{X}^Γ is known. The induced quantization error $\|X - \hat{X}^\Gamma\|_2$ is used to control the error (see below). These quantities related to the quantizer Γ are also called *companion parameters*.

Likewise, one can consider *a priori* the $\sigma(\hat{X}^\Gamma)$ -measurable random variable $F(\hat{X}^\Gamma)$ as a good approximation of the conditional expectation $\mathbb{E}(F(X) | \hat{X}^\Gamma)$.

4.1 Lipschitz functionals

Assume that the functional F is Lipschitz continuous on H . Then

$$\left| \mathbb{E}(F(X) | \hat{X}^\Gamma) - F(\hat{X}^\Gamma) \right| \leq [F]_{\text{Lip}} \mathbb{E}(|X - \hat{X}^\Gamma| | \hat{X}^\Gamma)$$

so that, for every real exponent $r \geq 1$,

$$\|\mathbb{E}(F(X) | \hat{X}^\Gamma) - F(\hat{X}^\Gamma)\|_r \leq [F]_{\text{Lip}} \|X - \hat{X}^\Gamma\|_r$$

(where we applied conditional Jensen inequality to the convex function $u \mapsto u^r$). In particular, using that $\mathbb{E} F(X) = \mathbb{E}(\mathbb{E}(F(X) | \hat{X}^\Gamma))$, one derives (with $r = 1$) that

$$\begin{aligned} \left| \mathbb{E} F(X) - \mathbb{E} F(\hat{X}^\Gamma) \right| &\leq \|\mathbb{E}(F(X) | \hat{X}^\Gamma) - F(\hat{X}^\Gamma)\|_1 \\ &\leq [F]_{\text{Lip}} \|X - \hat{X}^\Gamma\|_1. \end{aligned}$$

Finally, using the monotony of the $L^r(\mathbb{P})$ -norms as a function of r yields

$$\left| \mathbb{E} F(X) - \mathbb{E} F(\hat{X}^\Gamma) \right| \leq [F]_{\text{Lip}} \|X - \hat{X}^\Gamma\|_1 \leq [F]_{\text{Lip}} \|X - \hat{X}^\Gamma\|_2. \quad (5)$$

In fact, considering the Lipschitz functional $F(\xi) := d(\xi, \Gamma)$, shows that

$$\|X - \hat{X}^\Gamma\|_1 = \sup_{[F]_{\text{Lip}} \leq 1} \left| \mathbb{E} F(X) - \mathbb{E} F(\hat{X}^\Gamma) \right|. \quad (6)$$

The Lipschitz functionals making up a characterizing family for the weak convergence of probability measures on H , one derives that, for any sequence of N -quantizers Γ^N satisfying $\|X - \hat{X}^{\Gamma^N}\|_1 \rightarrow 0$ as $N \rightarrow \infty$,

$$\sum_{1 \leq i \leq N} \mathbb{P}(\hat{X}^{\Gamma^N} = x_i^N) \delta_{x_i^N} \xrightarrow{(H)} \mathbb{P}_X$$

where $\xrightarrow{(H)}$ denotes the weak convergence of probability measures on $(H, |\cdot|_H)$.

4.2 Differentiable functionals with Lipschitz differentials

Assume now that F is differentiable on H , with a Lipschitz continuous differential DF , and that the quantizer Γ is *stationary* (see Equation (4)).

A Taylor expansion yields

$$\left| F(X) - F(\hat{X}^\Gamma) - DF(\hat{X}^\Gamma) \cdot (X - \hat{X}^\Gamma) \right| \leq [DF]_{\text{Lip}} |X - \hat{X}^\Gamma|^2.$$

Taking conditional expectation given \hat{X}^Γ yields

$$\left| \mathbb{E}(F(X) | \hat{X}^\Gamma) - F(\hat{X}^\Gamma) - \mathbb{E} \left(DF(\hat{X}^\Gamma) \cdot (X - \hat{X}^\Gamma) | \hat{X}^\Gamma \right) \right| \leq [DF]_{\text{Lip}} \mathbb{E}(|X - \hat{X}^\Gamma|^2 | \hat{X}^\Gamma).$$

Now, using that the random variable $DF(\hat{X}^\Gamma)$ is $\sigma(\hat{X}^\Gamma)$ -measurable, one has

$$\mathbb{E} \left(DF(\hat{X}^\Gamma) \cdot (X - \hat{X}^\Gamma) \right) = \mathbb{E} \left(DF(\hat{X}^\Gamma) \cdot \mathbb{E}(X - \hat{X}^\Gamma | \hat{X}^\Gamma) \right) = 0$$

so that

$$\left| \mathbb{E}(F(X) | \hat{X}^\Gamma) - F(\hat{X}^\Gamma) \right| \leq [DF]_{\text{Lip}} \mathbb{E}(|X - \hat{X}^\Gamma|^2 | \hat{X}^\Gamma).$$

Then, for every real exponent $r \geq 1$,

$$\left\| \mathbb{E}(F(X) | \hat{X}^\Gamma) - F(\hat{X}^\Gamma) \right\|_r \leq [DF]_{\text{Lip}} \|X - \hat{X}^\Gamma\|_{2r}^2.$$

In particular, when $r = 1$, one derives like in the former setting

$$\left| \mathbb{E}F(X) - \mathbb{E}F(\hat{X}^\Gamma) \right| \leq [DF]_{\text{Lip}} \|X - \hat{X}^\Gamma\|_2^2. \quad (7)$$

In fact, the above inequality holds provided F is \mathcal{C}^1 with Lipschitz differential on every Voronoi cell $C_i(\Gamma)$. A similar characterization to (6) based on these functionals could be established.

Some variant of these cubature formulae can be found in [43] or [21] for functions or functionals F having only some local Lipschitz regularity.

4.3 Quantized approximation of $\mathbb{E}(F(X) | Y)$

Let X and Y be two H -valued random vector defined on the same probability space $(\Omega, \mathcal{A}, \mathbb{P})$ and $F : H \rightarrow \mathbb{R}$ be a Borel functional. The natural idea is to approximate $\mathbb{E}(F(X) | Y)$ by the quantized conditional expectation $\mathbb{E}(F(\hat{X}) | \hat{Y})$ where \hat{X} and \hat{Y} are quantizations of X and Y respectively.

Let $\varphi_F : H \rightarrow \mathbb{R}$ be a (Borel) version of the conditional expectation *i.e.* satisfying

$$\mathbb{E}(F(X) | Y) = \varphi_F(Y).$$

Usually, no closed form is available for the function φ_F but some regularity property can be established, especially in a (Feller) Markovian framework. Thus assume that both F and φ_F are Lipschitz continuous with Lipschitz coefficients $[F]_{\text{Lip}}$ and $[\varphi_F]_{\text{Lip}}$. Then

$$\mathbb{E}(F(X) | Y) - \mathbb{E}(F(\hat{X}) | \hat{Y}) = \mathbb{E}(F(X) | Y) - \mathbb{E}(F(X) | \hat{Y}) + \mathbb{E}(F(X) - F(\hat{X}) | \hat{Y}).$$

Hence, using that \hat{Y} is $\sigma(Y)$ -measurable and that conditional expectation is an L^2 -contraction,

$$\begin{aligned}
\|\mathbb{E}(F(X) | Y) - \mathbb{E}(F(X) | \hat{Y})\|_2 &= \|\mathbb{E}(F(X)|Y) - \mathbb{E}(\mathbb{E}(F(\hat{X})|Y)|\hat{Y})\|_2 \\
&\leq \|\varphi_F(Y) - \mathbb{E}(F(X)|\hat{Y})\|_2 \\
&= \|\varphi_F(Y) - \mathbb{E}(\varphi_F(Y)|\hat{Y})\|_2 \\
&\leq \|\varphi_F(Y) - \varphi_F(\hat{Y})\|_2.
\end{aligned}$$

The last inequality follows from the definition of conditional expectation given \hat{Y} as the best quadratic approximation among $\sigma(\hat{Y})$ -measurable random variables. On the other hand, still using that $\mathbb{E}(\cdot | \sigma(\hat{Y}))$ is an L^2 -contraction and this time that F is Lipschitz continuous yields

$$\|\mathbb{E}(F(X) - F(\hat{X}) | \hat{Y})\|_2 \leq \|F(X) - F(\hat{X})\|_2 \leq [F]_{\text{Lip}} \|X - \hat{X}\|_2.$$

Finally,

$$\|\mathbb{E}(F(X) | Y) - \mathbb{E}(F(\hat{X}) | \hat{Y})\|_2 \leq [F]_{\text{Lip}} \|X - \hat{X}\|_2 + [\varphi_F]_{\text{Lip}} \|Y - \hat{Y}\|_2.$$

In the non-quadratic case the above inequality remains valid provided $[\varphi_F]_{\text{Lip}}$ is replaced by $2[\varphi_F]_{\text{Lip}}$.

5 Vector quantization rate ($H = \mathbb{R}^d$)

The fact that $e_N(X, \mathbb{R}^d)$ is a non-increasing sequence that goes to 0 as N goes to ∞ is a rather simple result established in Proposition 1. Its rate of convergence to 0 is a much more challenging problem. An answer is provided by the so-called Zador Theorem stated below.

This theorem was first stated and established for distributions with compact supports by Zador (see [57, 58]). Then a first extension to general probability distributions on \mathbb{R}^d is developed in [5]. The first mathematically rigorous proof can be found in [17], and relies on a random quantization argument (Pierce Lemma).

Theorem 1. (a) SHARP RATE. *Let $r > 0$ and $X \in L^{r+\eta}(\mathbb{P})$ for some $\eta > 0$. Let $\mathbb{P}_X(d\xi) = \varphi(\xi) d\xi + \nu(d\xi)$ be the canonical decomposition of the distribution of X (ν and the Lebesgue measure are singular). Then (if $\varphi \not\equiv 0$),*

$$e_{N,r}(X, \mathbb{R}^d) \sim \tilde{J}_{r,d} \times \left(\int_{\mathbb{R}^d} \varphi^{\frac{d}{d+r}}(u) du \right)^{\frac{1}{d} + \frac{1}{r}} \times N^{-\frac{1}{d}} \quad \text{as } N \rightarrow +\infty. \quad (8)$$

where $\tilde{J}_{r,d} \in (0, \infty)$.

(b) NON ASYMPTOTIC UPPER BOUND (see e.g. [33]). *Let $d \geq 1$. There exists $C_{d,r,\eta} \in (0, \infty)$ such that, for every \mathbb{R}^d -valued random vector X ,*

$$\forall N \geq 1, \quad e_{N,r}(X, \mathbb{R}^d) \leq C_{d,r,\eta} \|X\|_{r+\eta} N^{-\frac{1}{d}}.$$

Remarks. • The real constant $\tilde{J}_{r,d}$ clearly corresponds to the case of the uniform distribution over the unit hypercube $[0,1]^d$ for which the slightly more precise statement holds

$$\lim_N N^{\frac{1}{d}} e_{N,r}(X, \mathbb{R}^d) = \inf_N N^{\frac{1}{d}} e_{N,r}(X, \mathbb{R}^d) = \tilde{J}_{r,d}.$$

The proof is based on a self-similarity argument. The value of $\tilde{J}_{r,d}$ depends on the reference norm on \mathbb{R}^d . When $d = 1$, elementary computations show that $\tilde{J}_{r,1} = (r+1)^{-\frac{1}{r}}/2$. When $d = 2$, with the canonical Euclidean norm, one shows (see [37] for a proof, see also [17]) that $\tilde{J}_{2,d} = \sqrt{\frac{5}{18\sqrt{3}}}$. Its exact value is unknown for $d \geq 3$ but, still for the canonical Euclidean norm, one has (see [17]) using some random quantization arguments,

$$\tilde{J}_{2,d} \sim \sqrt{\frac{d}{2\pi e}} \approx \sqrt{\frac{d}{17,08}} \quad \text{as } d \rightarrow +\infty.$$

- When $\varphi \equiv 0$ the distribution of X is purely singular. The rate (8) still holds in the sense that $\lim_N N^{\frac{1}{d}} e_{r,N}(X, \mathbb{R}^d) = 0$. Consequently, this is not the right asymptotics. The quantization problem for singular measures (like uniform distribution on fractal compact sets) has been extensively investigated by several authors, leading to the definition of a quantization dimension in connection with the rate of convergence of the quantization error on these sets. For more details we refer to [17, 18] and the references therein.

- A more naive way to quantize the uniform distribution on the unit hypercube is to proceed by *product quantization* i.e. by quantizing the marginals of the uniform distribution. If $N = m^d$, $m \geq 1$, one easily proves that the best quadratic product quantizer (for the canonical Euclidean norm on \mathbb{R}^d) is the “midpoint square grid”

$$\Gamma^{sq,N} = \left(\frac{2i_1 - 1}{2m}, \dots, \frac{2i_d - 1}{2m} \right)_{1 \leq i_1, \dots, i_d \leq m}$$

which induces a quadratic quantization error equal to

$$\sqrt{\frac{d}{12}} \times N^{-\frac{1}{d}}.$$

Consequently, product quantizers are still *rate optimal* in every dimension d . Moreover, note that the ratio of these two rates remains bounded as $d \uparrow \infty$.

6 Optimal quantization and QMC

The principle of Quasi-Monte Carlo method (QMC) is to approximate the integral of a function $f : [0,1]^d \rightarrow \mathbb{R}$ with respect to the uniform distribution

on $[0, 1]^d$, i.e. $\int_{[0, 1]^d} f d\lambda_d = \int_{[0, 1]^d} f(\xi^1, \dots, \xi^d) d\xi^1 \dots d\xi^d$ (λ_d denotes the Lebesgue measure on $[0, 1]^d$), by the uniformly weighted sum

$$\frac{1}{N} \sum_{k=1}^N f(x_k)$$

of values of f at the points of a so-called low discrepancy N -tuple (x_1, \dots, x_N) (or set). This N -tuple can be the first N terms of an infinite sequence.

If f has finite variations denoted $V(f)$ – either in the measure sense (see [4, 47]) or in the Hardy and Krause sense (see [38] p.19) – the Koksma-Hlawka inequality provides an upper bound for the integration error induced by this method, namely

$$\left| \frac{1}{N} \sum_{k=1}^N f(x_k) - \int_{[0, 1]^d} f d\lambda_d \right| \leq V(f) \text{Disc}_N^*(x_1, \dots, x_N)$$

where

$$\text{Disc}_N^*(x_1, \dots, x_N) := \sup_{y \in [0, 1]^d} \left| \frac{1}{N} \sum_{k=1}^N \mathbf{1}_{\{x_k \in \llbracket 0, y \rrbracket\}} - \lambda_d(\llbracket 0, y \rrbracket) \right|$$

(with $\llbracket 0, y \rrbracket = \prod_{i=1}^d [0, y^i]$, $y = (y^1, \dots, y^d) \in [0, 1]^d$). The error modulus $\text{Disc}_N^*(x_1, \dots, x_N)$ denotes the *discrepancy at the origin* of the N -tuple (x_1, \dots, x_N) . For every $N \geq 1$, there exists $[0, 1]^d$ -valued N -tuples $x^{(N)}$ such that

$$\text{Disc}_N^*(x^{(N)}) \leq C_d \frac{(\log N)^{d-1}}{N}, \quad (9)$$

where $C_d \in (0, \infty)$ is a real constant only depending on d . This result can be proved using the so-called Hammersely procedure (see e.g. [38], p.31). When $x^{(N)} = (x_1, \dots, x_N)$ is made of the first N terms of a $[0, 1]^d$ -valued sequence $(x_k)_{k \geq 1}$, then the above upper bound has to be replaced by $C'_d \frac{(\log N)^d}{N}$ ($C'_d \in (0, \infty)$). Such a sequence $x = (x_k)_{k \geq 1}$ is said to be a *sequence with low discrepancy* (see [38] and the references therein for a comprehensive theoretical overview, but also [4, 47] for examples supported by numerical tests). When one only has $\text{Disc}_N^*(x_1, \dots, x_N) \rightarrow 0$ as $N \rightarrow \infty$, the sequence is said to be *uniformly distributed in $[0, 1]^d$* .

It is widely shared by *QMC* specialists that these rates are (in some sense) optimal although this remains a conjecture except when $d = 1$. To be precise what is known and what is conjectured is the following:

– Any $[0, 1]^d$ -valued N -tuple $x^{(N)}$ satisfies $D_N^*(x^{(N)}) \geq B_d N^{-1} (\log N)^{\beta(d)}$ where $\beta(d) = \frac{d-1}{2}$ if $d \geq 2$ (see [52] and also [38] and the references therein), $\beta(1) = 0$ and $B_d > 0$ is a real constant only depending on d ; the conjecture is that $\beta(d) = d - 1$.

– Any $[0, 1]^d$ -valued sequence $(x_k)_{k \geq 1}$ satisfies $D_N^*(x^{(N)}) \geq B_d N^{-1} (\log N)^{\beta'(d)}$ for infinitely many N , where $\beta'(d) = \frac{d}{2}$ if $d \geq 2$ and $\beta'(1) = 1$ and $B'_d > 0$ is a real constant only depending on d ; the conjecture is that $\beta(d) = d$. This follows from the result for N -tuple by the Hammersley procedure (see *e.g.* [4]).

Furthermore, as concerns the use of Koksma-Hlawka inequality as an error bound for *QMC* numerical integration, the different notions of finite variation (which are closely connected) all become more and more restrictive – and subsequently less and less “natural” as a regularity property of functions – when the dimension d increases. Thus the Lipschitz continuous function f defined by $f(\xi^1, \xi^2, \xi^3) := (\xi^1 + \xi^2 + \xi^3) \wedge 1$ has infinite variation on $[0, 1]^3$.

When applying Quasi-Monte Carlo approximation of integrals with “standard” continuous functions on $[0, 1]^d$, the best known error bound, due to Proinov, is given by the following theorem.

Theorem 2. (*Proinov [50]*) (a) Assume \mathbb{R}^d is equipped with the ℓ^∞ -norm $|(u^1, \dots, u^d)|_\infty := \max_{1 \leq i \leq d} |u^i|$. Let $(x_1, \dots, x_N) \in ([0, 1]^d)^N$. For every continuous function $f : [0, 1]^d \rightarrow \mathbb{R}$,

$$\left| \int_{[0, 1]^d} f(u) du - \frac{1}{N} \sum_{k=1}^N f(x_k) \right| \leq K_d \omega_f((Disc_N^*(x_1, \dots, x_N))^{\frac{1}{d}})$$

where $\omega_f(\delta) := \sup_{x, y \in [0, 1]^d, |x-y|_\infty \leq \delta} |f(x) - f(y)|$, $\delta \in (0, 1)$, is the uniform continuity modulus of f (with respect to the ℓ_∞ -norm) and $C_d \in (0, \infty)$ is a universal constant only depending on d .

(b) If $d = 1$, $K_d = 1$ and if $d \geq 2$, $K_d \in [1, 4]$.

Remark. Note that if f is Lipschitz continuous, then $\omega_f(\delta) = [f]_{\text{Lip}} \delta$ where $[f]_{\text{Lip}}$ denotes the Lipschitz coefficient of f (with respect to the ℓ_∞ -norm).

First, this result emphasizes that low discrepancy sequences or sets do suffer from the *curse of dimensionality* when a *QMC* approximation is implemented on functions having a “natural” regularity like Lipschitz continuity.

One also derives from this theorem an inequality between $(L^1(\mathbb{P}), \ell_\infty)$ -quantization error of the uniform distribution $U([0, 1]^d)$ and the discrepancy at the origin of a N -tuple (x_1, \dots, x_N) , namely

$$\|U - \widehat{U}^{\{x_1, \dots, x_N\}}\|_{\ell^\infty} \leq K_d (Disc_N^*(x_1, \dots, x_N))^{\frac{1}{d}}$$

since the function $\xi \mapsto \min_{1 \leq k \leq N} |x_k - \xi|_\infty$ is clearly ℓ_∞ -Lipschitz continuous with Lipschitz coefficient 1. The inequality also follows from the characterization established in (6) (which is clearly still true for non Euclidean norms). Then one may derive some bounds for Euclidean norms (and in fact any norms) on \mathbb{R}^d (probably not sharp in terms of constant) since all the norms are strongly equivalent. However the bounds for optimal quantization

error derived from Zador's Theorem ($O(N^{-\frac{1}{d}})$) and those for low discrepancy sets (see (9)) suggest that overall, optimal quantization provides lower error bounds for numerical integration of Lipschitz functions than low discrepancy sets, at least for generic values of N . (However, standard computations show that for midpoint square grids (with $N = m^d$ points) both quantization errors and discrepancy behave like $\frac{1}{m} = N^{-\frac{1}{d}}$).

7 Optimal quadratic functional quantization of Gaussian processes

Optimal quadratic functional quantization of Gaussian processes is closely related to their so-called Karhunen-Loève expansion which can be seen in some sense as some infinite dimensional Principal Component Analysis (*PCA*) of a (Gaussian) process. Before stating a general result for Gaussian processes, we start by the standard Brownian motion: it is the most important example in view of (numerical) applications and for this process, everything can be made explicit.

7.1 Brownian motion

One considers the Hilbert space $H = L_T^2 := L^2([0, T], dt)$, $(f|g)_2 = \int_0^T f(t)g(t)dt$, $\|f\|_{L_T^2} = \sqrt{(f|f)_2}$. The covariance operator C_W of the Brownian motion $W = (W_t)_{t \in [0, T]}$ is defined on L_T^2 by

$$C_W(f) := \mathbb{E}((f, W)_2 W) = \left(t \mapsto \int_0^T (s \wedge t) f(s) ds \right).$$

It is a symmetric positive trace class operator which can be diagonalized in the so-called Karhunen-Loève (*K-L*) orthonormal basis $(e_n^W)_{n \geq 1}$ of L_T^2 , with eigenvalues $(\lambda_n)_{n \geq 1}$, given by

$$e_n^W(t) = \sqrt{\frac{2}{T}} \sin\left(\pi(n - \frac{1}{2})\frac{t}{T}\right), \quad \lambda_n = \left(\frac{T}{\pi(n - \frac{1}{2})}\right)^2, \quad n \geq 1.$$

This classical result can be established as a simple exercise by solving the functional equation $C_W(f) = \lambda f$. In particular, one can expand W itself on this basis so that

$$W \stackrel{L_T^2}{=} \sum_{n \geq 1} (W|e_n^W)_2 e_n^W.$$

Now, the orthonormality of the (*K-L*) basis implies, using Fubini's Theorem,

$$\mathbb{E}((W|e_k^W)_2(W|e_\ell^W)_2) = (e_k^W|C_W(e_\ell^W))_2 = \lambda_\ell \delta_{k\ell}$$

where $\delta_{k\ell}$ denotes the Kronecker symbol. Hence the Gaussian sequence $((W|e_n^W)_2)_{n \geq 1}$ is pairwise non-correlated which implies that these random variables are independent. The above identity also implies that $\text{Var}((W|e_n^W)_2) = \lambda_n$. Finally this shows that

$$W \stackrel{L_T^2}{=} \sum_{n \geq 1} \sqrt{\lambda_n} \xi_n e_n^W \quad (10)$$

where $\xi_n := (W|e_n^W)_2 / \sqrt{\lambda_n}$, $n \geq 1$, is an i.i.d. sequence of $\mathcal{N}(0; 1)$ -distributed random variables. Furthermore, this K - L expansion converges in a much stronger sense since $\sup_{t \in [0, T]} |W_t - \sum_{k=1}^n \sqrt{\lambda_k} \xi_k e_k^W(t)| \rightarrow 0$ \mathbb{P} -a.s. and

$$\left\| \sup_{[0, T]} |W_t - \sum_{1 \leq k \leq n} \sqrt{\lambda_k} \xi_k e_k^W(t)| \right\|_2 = O\left(\sqrt{\log n/n}\right)$$

(see *e.g.* [32]). Similar results (with various rates) hold true for a wide class of Gaussian processes expanded on “admissible” basis (see *e.g.* [34]).

Theorem 3. ([29] (2002) and [30] (2003)) *Let Γ^N , $N \geq 1$, be a sequence of optimal N -quantizers for W .*

(a) *For every $N \geq 1$, $\text{span}(\Gamma^N) = \text{span}\{e_1^W, \dots, e_{d(N)}^W\}$ with $d(N) = \Omega(\log N)$. Furthermore \widehat{W}^{Γ^N} and $W - \widehat{W}^{\Gamma^N}$ are independent.*

(b) $e_N(W, L_T^2) = \|W - \widehat{W}^{\Gamma^N}\|_2 \sim \frac{T\sqrt{2}}{\pi} \frac{1}{\sqrt{\log N}}$ as $N \rightarrow \infty$.

Remark. • The fact, confirmed by numerical experiments (see Section 7.3 Figure 7.3), that $d(N) \sim \log N$ holds as a conjecture.

• Denoting Π_d the orthogonal projection on $\text{span}\{e_1^W, \dots, e_d^W\}$, one derives from (a) that $\widehat{W}^{\Gamma^N} = \Pi_{d(N)}(W)^{\Gamma^N}$ (optimal quantization at level N) and

$$\begin{aligned} \|W - \widehat{W}^{\Gamma^N}\|_2^2 &= \|\Pi_{d(N)}(W) - \Pi_{d(N)}(W)^{\Gamma^N}\|_2^2 + \|W - \Pi_{d(N)}(W)\|_2^2 \\ &= e_N\left(Z_{d(N)}, \mathbb{R}^{d(N)}\right)^2 + \sum_{n \geq d(N)+1} \lambda_n \end{aligned}$$

where $Z_{d(N)} \stackrel{d}{=} \Pi_{d(N)}(W) \sim \bigotimes_{k=1}^{d(N)} \mathcal{N}(0; \lambda_k)$.

7.2 Centered Gaussian processes

The above Theorem 3 devoted to the standard Brownian motion is a particular case of a more general theorem which holds for a wide class of Gaussian processes

Theorem 4. ([29] (2002) and [30] (2004)) Let $X = (X_t)_{t \in [0, T]}$ be a Gaussian process with K -L eigensystem $(\lambda_n^X, e_n^X)_{n \geq 1}$ (with $\lambda_1 \geq \lambda_2 \geq \dots$ is non-increasing). Let Γ^N , $N \geq 1$, be a sequence of quadratic optimal N -quantizers for X . Assume

$$\lambda_n^X \sim \frac{\kappa}{n^b} \quad \text{as } n \rightarrow \infty \quad (b > 1).$$

(a) $\text{span}(\Gamma^N) = \text{span}\{e_1^X, \dots, e_{d^X(N)}^X\}$ and $d^X(N) = \Omega(\log N)$.

(b) $e_N(X, L_T^2) = \|X - \hat{X}^{\Gamma^N}\|_2 \sim \sqrt{\kappa} \sqrt{b^b(b-1)^{-1}} (2 \log N)^{-\frac{b-1}{2}}$.

Remarks. • The above result admits an extension to the case $\lambda_n^X \sim \varphi(n)$ as $n \rightarrow \infty$ with φ regularly varying, index $-b \leq -1$ (see [30]). In [29], upper or lower bounds are also established when

$$(\lambda_n^X \leq \varphi(n), \quad n \geq 1) \quad \text{or} \quad (\lambda_n^X \geq \varphi(n), \quad n \geq 1).$$

- The sharp asymptotics $d^X(N) \sim \frac{2}{b} \log N$ holds as a conjecture.

Applications to classical (centered) Gaussian processes.

- Brownian bridge: $X_t := W_t - \frac{t}{T} W_T$, $t \in [0, T]$ and $e_n^X(t) = \sqrt{2/T} \sin(\pi n \frac{t}{T})$, $\lambda_n = (\frac{T}{\pi n})^2$, so that $e_N(X, L_T^2) \sim T \frac{\sqrt{2}}{\pi} (\log N)^{-\frac{1}{2}}$.
- Fractional Brownian motion with Hurst constant $H \in (0, 1)$

$$e_N(W^H, L_T^2) \sim T^{H+\frac{1}{2}} c(H) (\log N)^{-H}$$

where $c(H) = \left(\frac{\Gamma(2H) \sin(\pi H) (1+2H)}{\pi} \right)^{\frac{1}{2}} \left(\frac{1+2H}{2\pi} \right)^H$ and $\Gamma(t)$ denotes the Gamma function at $t > 0$.

- Some further explicit sharp rates can be derived from the above theorem for other classes of Gaussian stochastic processes (see [30], 2004) like the fractional Ornstein-Uhlenbeck processes, the Gaussian diffusions, a wide class Gaussian stationary processes (the quantization rate is derived from the high frequency asymptotics of its spectral density, assumed to be square integrable on the real line), for the m -folded integrated Brownian motion, the fractional Brownian sheet, etc.

- Of course some upper bounds can be derived for some even wider classes of processes, based on the above first remark (see *e.g.* [29], 2002).

Extensions to $r, p \neq 2$ When the processes have some self-similarity properties, it is possible to obtain some sharp rates in the non purely quadratic case: this has been done for fractional Brownian motion in [12] using some quite different techniques in which self-similarity properties plays there a crucial role. It leads to the following sharp rates, for $p \in [1, +\infty]$ and $r \in (0, \infty)$

$$e_{N,r}(W^H, L_T^p) \sim T^{H+\frac{1}{2}} c(r, H) (\log N)^{-H}, \quad c(r, H) \in (0, +\infty).$$

7.3 Numerical optimization of quadratic functional quantization

Thanks to the scaling property of Brownian motion, one may focus on the normalized case $T = 1$. The numerical approach to optimal quantization of the Brownian motion is essentially based on Theorem 3 and the remark that follows: indeed these results show that quadratic optimal functional quantization of a centered Gaussian process reduces to a finite dimensional optimal quantization problem for a Gaussian distribution with a diagonal covariance structure. Namely the optimization problem at level N reads

$$(\mathcal{O}_N) \equiv \begin{cases} e_N(W, L_T^2) := e_N(Z_{d(N)}, \mathbb{R}^{d(N)})^2 + \sum_{k \geq d(N)+1} \lambda_k \\ \text{where } Z_{d(N)} \stackrel{d}{=} \bigotimes_{k=1}^{d(N)} \mathcal{N}(0, \lambda_k). \end{cases}$$

Moreover, if $\beta^N := \{\beta_1^N, \dots, \beta_N^N\}$ denotes an optimal N -quantizer of $Z_{d(N)}$, then, the optimal N -quantizer Γ^N of W reads $\Gamma^N = \{x_1^N, \dots, x_N^N\}$ with

$$x_i^N(t) = \sum_{1 \leq \ell \leq d(N)} (\beta_i^N)^\ell e_\ell^W(t), \quad i = 1, \dots, N. \quad (11)$$

The good news is that (\mathcal{O}_N) is in fact a *finite dimensional quantization optimization problem* for each $N \geq 1$. The bad news is that the problem is somewhat ill conditioned since the decrease of the eigenvalues of W is very steep for small values of n : $\lambda_1 = 0.40528 \dots$, $\lambda_2 = 0.04503 \dots \approx \lambda_1/10$. This is probably one reason for which former attempts to produce good quantization of the Brownian motion first focused on other kinds of quantizers like *scalar product quantizers* (see [44] and Section 7.4 below) or *d-dimensional block product quantizations* (see [56] and [35]).

Optimization of the (quadratic) quantization of \mathbb{R}^d -valued random vector has been extensively investigated since the early 1950's, first in 1-dimension, then in higher dimension when the cost of numerical Monte Carlo simulation was drastically cut down (see [15]). Recent application of optimal vector quantization to numerics turned out to be much more demanding in terms of accuracy. In that direction, one may cite [43], [36] (mainly focused on numerical optimization of the quadratic quantization of normal distributions). To apply the methods developed in these papers, it is more convenient to rewrite our optimization problem with respect to the standard d -dimensional distribution $\mathcal{N}(0; I_d)$ by simply considering the Euclidean norm derived from the covariance matrix $\text{Diag}(\lambda_1, \dots, \lambda_{d(N)})$ *i.e.*

$$(\mathcal{O}_N) \Leftrightarrow \begin{cases} N\text{-optimal quantization of } \bigotimes_{k=1}^{d(N)} \mathcal{N}(0, 1) \\ \text{for the covariance norm } |(z_1, \dots, z_{d(N)})|^2 = \sum_{k=1}^{d(N)} \lambda_k z_k^2. \end{cases}$$

The main point is of course that the dimension $d(N)$ is unknown. However (see Figure 7.3), one clearly verifies on small values of N that the conjecture ($d(N) \sim \log N$) is most likely true. Then for higher values of N one relies on it to shift from one dimension to another following the rule $d(N) = d$, $N \in \{e^d, \dots, e^{d+1} - 1\}$.

A toolbox for quantization optimization: a short overview

Here is a short overview of stochastic optimization methods to compute optimal or at least locally optimal quantizers in finite dimension. For more details we refer to [43] and the references therein. Let $Z \stackrel{d}{=} \mathcal{N}(0; I_d)$.

Competitive Learning Vector Quantization (CLVQ). This procedure is a recursive stochastic approximation gradient descent based on the integral representation of the gradient $\nabla D_N^Z(x)$, $x \in H^n$ (temporarily coming back to N -tuple notation) of the distortion as the expectation of a *local gradient i.e.*

$$\forall x^N \in H^N, \quad \nabla D_N^Z(x^N) = \mathbb{E}(\nabla D_N^Z(x^N, \zeta)), \quad \zeta_k \text{ i.i.d.}, \quad \zeta_1 \stackrel{d}{=} \mathcal{N}(0, I_d)$$

so that, starting from $x^N(0) \in (\mathbb{R}^d)^N$, one sets

$$\forall k \geq 0, \quad x^N(k+1) = x^N(k) - \frac{c}{k+1} \nabla D_N^Z(x^N(k), \zeta_{k+1})$$

where $c \in (0, 1]$ is a real constant to be tuned. As set, this looks quite formal but the operating *CLVQ* procedure consists of two phases at each iteration:

(i) *Competitive Phase*: Search of the nearest neighbor $x^N(k)_{i^*(k+1)}$ of ζ_{k+1} among the components of $x^N(k)_i$, $i = 1, \dots, N$ (using a “winning convention” in case of conflict on the boundary of the Voronoi cells).

(ii) *Cooperative Phase*: One moves the winning component toward ζ_{k+1} using a dilatation i.e. $x^N(k+1)_{i^*(k+1)} = \text{Dilatation}_{\zeta_{k+1}, 1 - \frac{c}{k+1}}(x^N(k)_{i^*(k+1)})$.

This procedure is useful for small or medium values of N . For an extensive study of this procedure, which turns out to be singular in the world of recursive stochastic approximation algorithms, we refer to [40]. For general background on stochastic approximation, we refer to [25, 3].

The randomized “Lloyd I procedure”. This is the randomization of the stationarity based fixed point procedure since any optimal quantizer satisfies (4):

$$\widehat{Z}^{x^N(k+1)} = \mathbb{E}(Z \mid \widehat{Z}^{x^N(k)}), \quad x^N(0) \subset \mathbb{R}^d.$$

At every iteration the conditional expectation $\mathbb{E}(Z \mid \widehat{Z}^{x^N(k)})$ is computed using a Monte Carlo simulation. For more details about practical aspects of Lloyd I procedure we refer to [43]. In [36], an approach based on genetic evolutionary algorithms is developed.

For both procedures, one may substitute a sequence of quasi-random numbers to the usual pseudo-random sequence. This often speeds up the rate of

convergence of the method, although this can only be proved (see [27]) for a very specific class of stochastic algorithm (to which *CLVQ* does not belong).

The most important step to preserve the accuracy of the quantization as N (and $d(N)$) increase is to use the so-called *splitting method* which finds its origin in the proof of the existence of an optimal N -quantizer: once the optimization of a quantization grid of size N is achieved, one specifies the starting grid for the size $N + 1$ or more generally $N + \nu$, $\nu \geq 1$, by merging the optimized grid of size N resulting from the former procedure with ν points sampled independently from the normal distribution with probability density proportional to $\varphi^{\frac{d}{d+2}}$ where φ denotes the p.d.f. of $\mathcal{N}(0; I_d)$. This rather unexpected choice is motivated by the fact that this distribution provides the lowest *in average* random quantization error (see [6]).

As a result, to be downloaded on the website [45] devoted to quantization:

www.quantize.maths-fi.com

◦ *Optimized stationary codebooks for W* : in practice, the N -quantizers β^N of the distribution $\otimes_{k=1}^{d(N)} \mathcal{N}(0; \lambda_k)$, $N = 1$ up to 10 000 ($d(N)$ runs from 1 up to 9).

◦ *Companion parameters*:

- distribution of \widehat{W}^{Γ^N} : $\mathbb{P}(\widehat{W}^{\Gamma^N} = x_i^N) = \mathbb{P}(\widehat{Z}_{d(N)}^{\beta^N} = \beta_i^N)$ (\leftarrow in $\mathbb{R}^{d(N)}$).
- The quadratic quantization error: $\|W - \widehat{W}^{\Gamma^N}\|_2$.

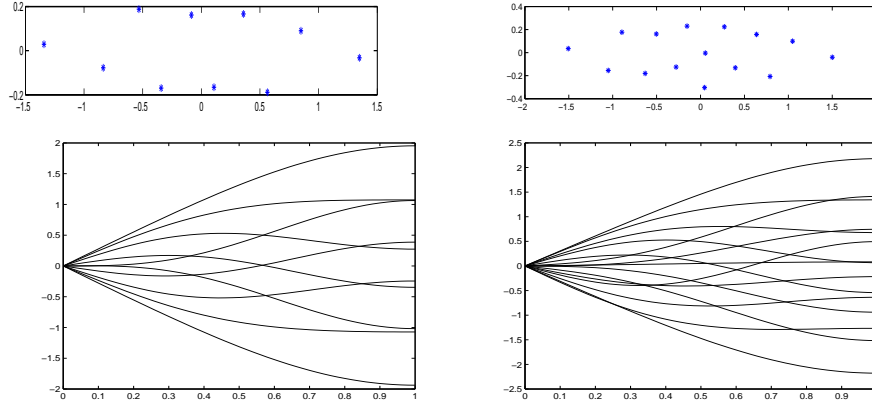


Fig. 3. Optimized functional quantization of the Brownian motion W for $N = 10, 15$ ($d(N) = 2$). Top: β^N depicted in \mathbb{R}^2 . Bottom: the optimized N -quantizer Γ^N .




Fig. 4. Optimized functional quantization of the Brownian motion W . The N -quantizers Γ^N . Left: $N = 48$ ($d(N) = 3$). Right: $N = 96$, $d(96) = 4$.




Fig. 5. Optimized N -quantizer Γ^N of the Brownian motion W with $N = 400$. The grey level of the paths codes their weights.




Fig. 6. Optimal functional quantization of the Brownian motion. $N \mapsto \log N (e_N(W, L_T^2))^2$, $N \in \{6, \dots, 160\}$. Vertical dashed lines: critical dimensions for $d(N)$, $e^2 \approx 7$, $e^3 \approx 20$, $e^4 \approx 55$, $e^5 \approx 148$.

7.4 An alternative: product functional quantization

Scalar Product functional quantization is a quantization method which produces rate optimal sub-optimal quantizers. They were used *e.g.* in [29] to provide exact rate (although not sharp) for a very large class of processes. The first attempts to use functional quantization for numerical computation with the Brownian motion was achieved with these quantizers (see [44]). We will see further on their assets. What follows is presented for the Brownian motion but would work for a large class of centered Gaussian processes.

Let us consider again the expansion of W in its K - L basis :

$$W \stackrel{L^2}{=} \sum_{n \geq 1} \sqrt{\lambda_n} \xi_n e_n^W$$

where $(\xi_n)_{n \geq 1}$ is an i.i.d. sequence $\mathcal{N}(0; 1)$ -distributed random variables (keep in mind this convergence also holds *a.s.* uniformly in $t \in [0, T]$). The idea is simply to quantize these (normalized) random coordinates ξ_n : for every $n \geq 1$, one considers an optimal N_n -quantization of ξ_n , denoted $\hat{\xi}_n^{(N_n)}$ ($N_n \geq 1$). For $n > m$, set $N_n = 1$ and $\hat{\xi}_n^{(N_n)} = 0$ (which is the optimal 1-quantization). The integer m is called the *length* of the product quantization. Then, one sets

$$\widehat{W}_t^{(N_1, \dots, N_m, prod)} := \sum_{n \geq 1} \sqrt{\lambda_n} \hat{\xi}_n^{(N_n)} e_n^W(t) = \sum_{n=1}^m \sqrt{\lambda_n} \hat{\xi}_n^{(N_n)} e_n^W(t).$$

Such a quantizer takes $\prod_{n=1}^m N_n \leq N$ values.

If one denotes by $\alpha^M = \{\alpha_1^M, \dots, \alpha_M^M\}$ the (unique) optimal quadratic M -quantizer of the $\mathcal{N}(0; 1)$ -distribution, the underlying quantizer of the above quantization $\widehat{W}^{(N_1, \dots, N_m, prod)}$ can be expressed as follows (if one introduces the appropriate multi-indexation): for every multi-index $\underline{i} := (i_1, \dots, i_m) \in \prod_{n=1}^m \{1, \dots, N_n\}$, set

$$x_{\underline{i}}^{(N)}(t) := \sum_{n=1}^m \sqrt{\lambda_n} \alpha_{i_n}^{(N_n)} e_n^W(t) \text{ and } \Gamma^{N_1, \dots, N_m, prod} := \left\{ x_{\underline{i}}^{(N)}, \underline{i} \in \prod_{n=1}^m \{1, \dots, N_n\} \right\}.$$

Then the product quantization $\widehat{W}^{(N_1, \dots, N_m, prod)}$ can be rewritten as

$$\widehat{W}_t^{(N_1, \dots, N_m, prod)} = \sum_{\underline{i}} \mathbf{1}_{\{W \in C_{\underline{i}}(\Gamma^{N_1, \dots, N_m, prod})\}} x_{\underline{i}}^{(N)}(t).$$

where the Voronoi cell of $x_{\underline{i}}^{(N)}$ is given by

$$C_{\underline{i}}(\Gamma^{N_1, \dots, N_m, prod}) = \prod_{n=1}^m (\alpha_{i_n - \frac{1}{2}}^{(N_n)}, \alpha_{i_n + \frac{1}{2}}^{(N_n)})$$

with $\alpha_{i \pm \frac{1}{2}}^{(M)} := \frac{\alpha_i^{(M)} + \alpha_{i \pm 1}^{(M)}}{2}$, $\alpha_0 = -\infty$, $\alpha_{M+1} = +\infty$.

Quantization rate by product quantizers

It is clear that the optimal product quantizer is the solution to the optimal integral bit allocation

$$\min \left\{ \|W - \widehat{W}^{(N_1, \dots, N_m, prod)}\|_2, N_1, \dots, N_m \geq 1, N_1 \times \dots \times N_m \leq N, m \geq 1 \right\}. \quad (12)$$

Expanding $\|W - \widehat{W}^{(N_1, \dots, N_m, prod)}\|_2^2 = \| |W - \widehat{W}^{(N_1, \dots, N_m, prod)}|_{L_T^2} \|_2^2$ yields

$$\|W - \widehat{W}^{(N_1, \dots, N_m, prod)}\|_2^2 = \sum_{n \geq 1} \lambda_n \|\widehat{\xi}_n^{(N_n)} - \xi_n\|_2^2 \quad (13)$$

$$= \sum_{n=1}^m \lambda_n (e_{N_n}^2(\mathcal{N}(0; 1), \mathbb{R}) - 1) + \frac{T^2}{2} \quad (14)$$

since
$$\sum_{n \geq 1} \lambda_n = \mathbb{E} \sum_{n \geq 1} (W | e_n^W)_2^2 = \mathbb{E} \int_0^T W_t^2 dt = \int_0^T t dt = \frac{T^2}{2}.$$

Theorem 5. (see [29]) *For every $N \geq 1$, there exists an optimal scalar product quantizer of size at most N (or at level N), denoted $\widehat{W}^{(N, prod)}$, of the Brownian motion defined as the solution to the minimization problem (12). Furthermore these optimal product quantizers make up a rate optimal sequence: there exists a real constant $c_W > 0$ such that*

$$\|W - \widehat{W}^{(N, prod)}\|_2 \leq \frac{c_W T}{(\log N)^{\frac{1}{2}}}.$$

Proof (sketch of). By scaling one may assume without loss of generality that $T = 1$. Combining (13) and Zador's Theorem shows

$$\begin{aligned} \|W - \widehat{W}^{(N_1, \dots, N_m, prod)}\|_2^2 &\leq C \left(\sum_{n=1}^m \frac{1}{n^2 N_n^2} \right) + \sum_{n \geq m+1} \lambda_n \\ &\leq C' \left(\sum_{n=1}^m \frac{1}{n^2 N_n^2} + \frac{1}{m} \right) \end{aligned}$$

with $\prod_n N_n \leq N$. Setting $m := m(N) = \lfloor \log N \rfloor$ and $N_k = \left\lceil \frac{(m!N)^{\frac{1}{m}}}{k} \right\rceil \geq 1$, $k = 1, \dots, m$, yields the announced upper-bound. \diamond

Remarks. • One can show that the length $m(N)$ of the optimal quadratic product quantizer satisfies

$$m(N) \sim \log N \quad \text{as} \quad N \rightarrow +\infty.$$

- The most striking fact is that very few ingredients are necessary to make the proof work as far as the quantization rate is concerned. We only need the basis of L_T^2 on which W is expanded to be orthonormal *or* the random coordinates to be orthogonal in $L^2(\mathbb{P})$. This robustness of the proof has been used to obtain some upper bounds for very wide classes of Gaussian processes by considering alternative orthonormal basis of L_T^2 like the Haar basis for processes having self-similarity properties (see [29]), or trigonometric basis for stationary processes (see [29]). More recently, combined with the non asymptotic Zador's Theorem, it was used to provide some connections between mean regularity of stochastic processes and quantization rate (see Section 10 and [33]).
- Block quantizers combined with large deviations estimates were used to provide the sharp rate obtained in Theorem 3 in [30].
- d -dimensional block quantization is also possible, possibly with varying block size, providing a constructive approach to sharp rate, see [56] and [35].
- A similar approach can also provide some $L^r(\mathbb{P})$ -rates for product quantization with respect to the sup-norm over $[0, T]$, see [32].

How to use product quantizers for numerical computations ?

For numerics one can assume by a scaling argument that $T = 1$. To use product quantizers for numerics we need to have access to the quantizers (or grid) at a given level N , their weights (and the quantization error). All these quantities are available with product quantizers. In fact the first attempts to use functional quantization for numerics (path dependent option pricing) were carried out with product quantizers (see [44]).

- The optimal product quantizers (denoted $I^{(N, prod)}$) at level N are explicit, given the optimal quantizers of the scalar normal distribution $\mathcal{N}(0; 1)$. In fact the optimal allocation of the size N_i of each marginal has been already achieved up to very high values of N . Some typical optimal allocation (and the resulting quadratic quantization error) are reported in the table below.

N	N_{rec}	Quant. Error	Opti. Alloc.
1	1	0.7071	1
10	10	0.3138	5-2
100	96	0.2264	12-4-2
1 000	966	0.1881	23-7-3-2
10 000	9 984	0.1626	26-8-4-3-2-2
100 000	97 920	0.1461	34 - 10 - 6 - 4 - 3 - 2 - 2

- The weights $\mathbb{P}(\widehat{W}^{(N, prod)} = x_{\underline{i}})$ are explicit too: the normalized coordinates ξ_n of W in its K - L basis are independent, consequently

$$\mathbb{P}(\widehat{W}^{(N, prod)} = x_{\underline{i}}) = \mathbb{P}(\widehat{\xi}_n^{(N_n)} = \alpha_{i_n}^{(N_n)}, n = 1, \dots, m(N))$$

$$= \prod_{n=1}^{m(N)} \underbrace{\mathbb{P}(\widehat{\xi}_n^{(N_n)} = \alpha_{i_n}^{(N_n)})}_{1D \text{ (tabulated) weights}}.$$

• Equation (14) shows that the (squared) quantization error of a product quantizer can be straightforwardly computed as soon as one knows the eigenvalues and the (squared) quantization error of the normal distributions for the N_i 's.

The optimal allocations up to $N = 12\,000$ can be downloaded on the website [45] as well as the necessary 1-dimensional optimal quantizers (including the weights and the quantization error) of the scalar normal distribution (up to a size of 500 which is quite enough for this purpose).

For numerical purpose we are also interested in the stationarity property since such quantizers produce lower (weak) errors in cubature formulas.

Proposition 2. (see [44]) *The product quantizers obtained from the K-L basis are stationary quantizers (although sub-optimal).*

Proof. Firstly, note that

$$\widehat{W}^{N,prod} = \sum_{n \geq 1} \sqrt{\lambda_n} \widehat{\xi}_n^{(N_n)} e_n(t)$$

so that $\sigma(\widehat{W}^{N,prod}) = \sigma(\widehat{\xi}_k^{(N_k)}, k \geq 1)$. Consequently

$$\begin{aligned} \mathbb{E}(W | \widehat{W}^{N,prod}) &= \mathbb{E}(W | \sigma(\widehat{\xi}_k^{(N_k)}, k \geq 1)) \\ \mathbb{E}(W | \widehat{W}^{N,prod}) &= \sum_{n \geq 1} \sqrt{\lambda_n} \mathbb{E}(\xi_n | \sigma(\widehat{\xi}_k^{(N_k)}, k \geq 1)) e_n^W \\ &\stackrel{i.i.d.}{=} \sum_{n \geq 1} \sqrt{\lambda_n} \mathbb{E}(\xi_n | \widehat{\xi}_n^{(N_n)}) e_n^W \\ &= \sum_{n \geq 1} \sqrt{\lambda_n} \widehat{\xi}_n^{(N_n)} e_n^W = \widehat{W}. \quad \diamond \end{aligned}$$

Remarks. • This result is no longer true for product quantizers based on other orthonormal basis.

• This shows the existence of non optimal stationary quantizers.

7.5 Optimal vs product quadratic functional quantization ($T = 1$)

◦ (NUMERICAL) OPTIMIZED QUANTIZATION: By scaling, we can assume without loss of generality that $T = 1$. We carried out a huge optimization task in order to produce some *optimized* quantization grids for the Brownian motion by solving numerically (\mathcal{O}_N) for $N = 1$ up to $N = 10\,000$.

$$e_N(W, L_T^2)^2 \approx \frac{0.2195}{\log N}, \quad N = 1, \dots, 10\,000.$$




Fig. 7. Product quantization of the Brownian motion: the N_{rec} -quantizer $I^{(N, \text{prod})}$, $N = 10$: $N_{\text{rec}} = 10$ and $N = 50$: $N_{\text{rec}} = 12 \times 4 = 48$.




Fig. 8. Product quantization of the Brownian motion: the N_{rec} -quantizer $I^{(N, \text{prod})}$, $N = 100$: $N_{\text{rec}} = 12 \times 4 \times 2 = 96$.

This value (see Figure 9(left)) is significantly greater than the theoretical (asymptotic) bound given by Theorem 3 which is

$$\lim_N \log Ne_N(W, L_T^2)^2 = \frac{2}{\pi^2} = 0.2026...$$

Our guess, supported by our numerical experiments, is that in fact $N \mapsto \log Ne_N(W, L_T^2)^2$ is possibly not monotonous but unimodal.

◦ OPTIMAL PRODUCT QUANTIZATION: as displayed on Figure 9(right), one has approximately

$$\min \left\{ \| |W - \widehat{W}|_{L_T^2} \|^2, 1 \leq N_1 \cdots N_m \leq N, m \geq 1 \right\} = \| W - \widehat{W}^{(N, \text{prod})} \|^2 \approx \frac{0.245}{\log N}$$

◦ OPTIMAL d -DIMENSIONAL BLOCK PRODUCT QUANTIZATION: let us briefly mention this approach developed in [56] in which product quantization is achieved by quantizing some marginal blocks of size 1, 2 or 3. By this approach, the corresponding constant is approximately 0.23, *i.e.* roughly in between scalar product quantization and optimized numeric quantization.

The conclusion, confirmed by our numerical experiments on option pricing (see Section 9), is that

- Optimal quantization is significantly more accurate on numerical experiments but is much more demanding since it needs to keep off line or at least to handle large files (say 1 GB for $N = 10\,000$).
- Both approaches are included in the option pricer PREMIA (MATHFI Project, Inria). An online benchmark is available on the website [45].




Fig. 9. Numerical quantization rates. Top (Optimal quantization). Line +++: $\log N \mapsto (\|W - \widehat{W}^N\|_2)^{-2}$. Dashed line: $\log N \mapsto \log N/0.2194$. Solid line: $\log N \mapsto \log N/0.25$. Bottom (Product quantization). Line +++: $\log N \mapsto (\min_{1 \leq k \leq N} \|W - \widehat{W}^{k,prod}\|_2^2)^{-1}$. Solid line: $\log N \mapsto \log N/0.25$.

8 Constructive functional quantization of diffusions

8.1 Rate optimality for Scalar Brownian diffusions

One considers on a probability space $(\Omega, \mathcal{A}, \mathbb{P})$ an homogenous Brownian diffusion process:

$$dX_t = b(X_t)dt + \vartheta(X_t)dW_t, \quad X_0 = x_0 \in \mathbb{R},$$

where b and ϑ are continuous on \mathbb{R} with at most linear growth (i.e. $|b(x)| + |\sigma(x)| \leq C(1 + |x|)$) so that at least a weak solution to the equation exists.

To devise a constructive way to quantize the diffusion X , it seems natural to start from a rate optimal quantization of the Brownian motion and to obtain some “good” (but how good?) quantizers for the diffusion by solving an appropriate *ODE*. So let $\Gamma^N = (w_1^N, \dots, w_N^N)$, $N \geq 1$, be a sequence of stationary rate optimal N -quantizers of W . One considers the following (non-coupled) Integral Equations:

$$dx_i^{(N)}(t) = \left(b(x_i^{(N)}(t)) - \frac{1}{2} \vartheta \theta'(x_i^{(N)}(t)) \right) dt + \vartheta(t, x_i^{(N)}(t)) dw_i^N(t). \quad (15)$$

Set

$$\tilde{X}_t^{x^{(N)}} = \sum_{k=1}^N x_i^{(N)}(t) \mathbf{1}_{\{\widehat{W}^{\Gamma^N} = w_i^N\}}.$$

The process $\tilde{X}^{x^{(N)}}$ is a *non-Voronoi* quantizer (since it is defined using the Voronoi diagram of W). What is interesting is that it is a *computable quantizer* (once the above integral equations have been solved) since the weights $\mathbb{P}(\widehat{W}^{\Gamma^N} = w_i^N)$ are known. The Voronoi quantization defined by $x^{(N)}$ induces a lower quantization error but we have no access to its weights for numerics. The good news is that $\tilde{X}^{x^{(N)}}$ is already rate optimal.

Theorem 6. ([31] (2006)) *Assume that b is differentiable, ϑ is positive twice differentiable and that $b' - b \frac{\vartheta'}{\vartheta} - \frac{1}{2} \vartheta \vartheta''$ is bounded. Then*

$$e_N(X, L_T^2) \leq \|X - \tilde{X}^{x^{(N)}}\|_2 = O((\log N)^{-\frac{1}{2}}).$$

If furthermore, $\vartheta \geq \varepsilon_0 > 0$, then $e_N(X, L_T^2) \approx (\log N)^{-\frac{1}{2}}$.

Remarks. • For some results in the non homogenous case, we refer to [31]. Furthermore, the above estimates still hold true for the $(L^r(\mathbb{P}), L_T^p)$ -quantization, $1 < r, p < +\infty$ provided $\|W - \widehat{W}^{\Gamma^N}\|_{L_T^p} = O((\log N)^{-\frac{1}{2}})$.

• This result is closely connected to the Doss-Sussman approach (see *e.g.* [13]) and in fact the results can be extended to some classes multi-dimensional diffusions (whose diffusion coefficient is the inverse of the gradient of a diffeomorphism) which include several standard multi-dimensional financial models (including the Black-Scholes model).

• A sharp quantization rate $e_{N,r}(X, L_T^p) \sim c(\log N)^{-\frac{1}{2}}$ for scalar elliptic diffusions is established in [10, 11] using a non constructive approach, $1 \leq p \leq \infty$.

EXAMPLE: Rate optimal product quantization of the Ornstein-Uhlenbeck process.

$$dX_t = -kX_t dt + \vartheta dW_t, \quad X_0 = x_0.$$

One solves the non-coupled integral (linear) system

$$x_i(t) = x_0 - k \int_0^t x_i(s) ds + \vartheta w_i^N(t),$$

where $\Gamma^N := \{w_1^N, \dots, w_N^N\}$, $N \geq 1$ is a *rate optimal* sequence of quantizers

$$w_i^N(t) = \sqrt{\frac{2}{T}} \sum_{\ell \geq 1} \varpi_{i,\ell} \frac{T}{\pi(\ell - 1/2)} \sin\left(\pi(\ell - 1/2) \frac{t}{T}\right), \quad i \in I_N.$$

If Γ^N is optimal for W then $\varpi_{i,\ell} := (\beta_i^N)^\ell$, $i = 1, \dots, N$, $1 \leq \ell \leq d(N)$ with the notations introduced in (11). If Γ^N is an optimal product quantizer (and $N_1, \dots, N_\ell, \dots$ denote the optimal size allocation), then $\varpi_{i,\ell} = \alpha_{i_\ell}^{(N_\ell)}$, where $i := (i_1, \dots, i_\ell, \dots) \in \prod_{\ell \geq 1} \{1, \dots, N_\ell\}$. Elementary computations show that

$$x_i^N(t) = e^{-kt} x_0 + \vartheta \sum_{\ell \geq 1} \chi_{i_\ell}^{(N_\ell)} \tilde{c}_\ell \varphi_\ell(t)$$

$$\text{with} \quad \tilde{c}_\ell = \frac{T^2}{(\pi(\ell - 1/2))^2 + (kT)^2}$$

$$\text{and} \quad \varphi_\ell(t) := \sqrt{\frac{2}{T}} \left(\frac{\pi}{T} (\ell - 1/2) \sin\left(\pi(\ell - 1/2) \frac{t}{T}\right) + k \left(\cos\left(\pi(\ell - 1/2) \frac{t}{T}\right) - e^{-kt} \right) \right).$$

8.2 Multi-dimensional diffusions for Stratanovich SDE's

The correcting term $-\frac{1}{2}\vartheta\vartheta'$ coming up in the integral equations suggest to consider directly some diffusion in the Stratanovich sense

$$dX_t = b(t, X_t) dt + \vartheta(t, X_t) \circ dW_t \quad X_0 = x_0 \in \mathbb{R}^d, \quad t \in [0, T].$$

(see e.g. [51] for an introduction) where $W = (W^1, \dots, W^d)$ is a d -dimensional standard Brownian Motion.

In that framework, we need to introduce the notion of p -variation: a continuous function $x : [0, T] \rightarrow \mathbb{R}^d$ has finite p -variations if

$$Var_{p,[0,T]}(x) := \sup \left\{ \left(\sum_{i=0}^{k-1} |x(t_i) - x(t_{i+1})|^p \right)^{\frac{1}{p}}, 0 \leq t_0 \leq t_1 \leq \dots \leq t_k \leq T, k \geq 1 \right\} < +\infty.$$

Then $d_p(x, x') = |x(0) - x'(0)| + Var_{p,[0,T]}(x - x')$ defines a distance on the set of functions with finite p -variations. It is classical background that $Var_{p,[0,T]}(W(\omega)) < +\infty$ $\mathbb{P}(d\omega)$ -a.s. for every $p > 2$.

One way to quantize W at level (at most) N is to quantize each component W^i at level $\lfloor \sqrt[p]{N} \rfloor$. One shows (see [30]) that $\|W - (\widehat{W}^1, \lfloor \sqrt[p]{N} \rfloor, \dots, \widehat{W}^d, \lfloor \sqrt[p]{N} \rfloor)\|_2 = O((\log N)^{-\frac{1}{2}})$.

Let $\mathcal{C}_b^r([0, T] \times \mathbb{R}^d)$ $r > 0$, denote the set of $\lfloor r \rfloor$ -times differentiable bounded functions $f : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ with bounded partial derivatives up to order $\lfloor r \rfloor$ and whose partial derivatives of order $\lfloor r \rfloor$ are $(r - \lfloor r \rfloor)$ -Hölder.

Theorem 7. (see [46]) Let $b, \vartheta \in \mathcal{C}_b^{2+\alpha}([0, T] \times \mathbb{R}^d)$ ($\alpha > 0$) and let $\Gamma^N = \{w_1^N, \dots, w_N^N\}$, $N \geq 1$, be a sequence of N -quantizers of the standard d -dimensional Brownian motion W such that $\|W - \widehat{W}^{\Gamma^N}\|_2 \rightarrow 0$ as $N \rightarrow \infty$. Let

$$\widetilde{X}_t^{x^{(N)}} := \sum_{i=1}^N x_i^{(N)}(t) \mathbf{1}_{\{\widehat{W}=w_i^N\}}$$

where, for every $i \in \{1, \dots, N\}$, $x_i^{(N)}$ is solution to

$$ODE_i \quad \equiv \quad dx_i^{(N)}(t) = b(t, x_i^{(N)}(t))dt + \vartheta(t, x_i^{(N)}(t))dw_i^N(t), \quad x_i^{(N)}(0) = x.$$

Then, for every $p \in (2, \infty)$,

$$Var_{p, [0, T]}(\widetilde{X}^{x^{(N)}} - X) \xrightarrow{\mathbb{P}} 0 \quad \text{as } N \rightarrow \infty.$$

Remarks. • The keys of this results are the Kolmogorov criterion, stationarity (in a slightly extended sense) and the connection with rough paths theory (see [28] for an introduction to rough paths theory, convergence in p -variation, etc).

• In that general setting we have no convergence rate although we conjecture that $\widetilde{X}^{x^{(N)}}$ remains rate optimal if \widehat{W}^{Γ^N} is.

• There are also some results about the convergence of stochastic integrals of the form $\int_0^t g(\widehat{W}_s^N) d\widehat{B}_s^N \rightarrow \int_0^t g(W_s) \circ dB_s$, with some rates of convergence when $W = B$ or W and B independent (depending on the regularity of the function g , see [46]).

9 Applications to path-dependent option pricing

The typical functionals F defined on $(L_T^2, |\cdot|_{L_T^2})$ for which $\mathbb{E}(F(W))$ can be approximated by the cubature formulae (5), (7) are of the form $F(\omega) := \varphi\left(\int_0^T f(t, \omega(t))dt\right) \mathbf{1}_{\{\omega \in \mathcal{C}([0, T], \mathbb{R})\}}$ where $f : [0, T] \times \mathbb{R} \rightarrow \mathbb{R}$ is *locally Lipschitz continuous* in the second variable, namely

$$\forall t \in [0, T], \forall u, v \in \mathbb{R}, |f(t, u) - f(t, v)| \leq C_f |u - v| (1 + g(|u|) + g(|v|))$$

(with $g : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is increasing, convex and $g(\sup_{t \in [0, T]} |W_t|) \in L^2(\mathbb{P})$) and $\varphi : \mathbb{R} \rightarrow \mathbb{R}$ is Lipschitz continuous. One could consider for ω some càdlàg functions as well. A classical example is the Asian payoff in a Black-Scholes model

$$F(\omega) = \exp(-rT) \left(\frac{1}{T} \int_0^T s_0 \exp(\sigma \omega(t) + (r - \sigma^2/2)t) dt - K \right)_+.$$

9.1 Numerical integration (II): log-Romberg extrapolation

Let $F : L_T^2 \rightarrow \mathbb{R}$ be a 3 times $|\cdot|_{L_T^2}$ -differentiable functional with bounded differentials. Assume $\widehat{W}^{(N)}$, $N \geq 1$, is a sequence of a rate-optimal stationary quantizations of the standard Brownian motion W . Assume furthermore that

$$\mathbb{E} \left(D^2 F(\widehat{W}^{(N)}) \cdot (W - \widehat{W}^{(N)})^{\otimes 2} \right) \sim \frac{c}{\log N} \quad \text{as } N \rightarrow \infty \quad (16)$$

and

$$\mathbb{E} |W - \widehat{W}^{(N)}|_{L_T^2}^3 = O \left((\log N)^{-\frac{3}{2}} \right). \quad (17)$$

Then, a higher order Taylor expansion yields

$$\begin{aligned} F(W) &= F(\widehat{W}^{(N)}) + DF(\widehat{W}^{(N)}) \cdot (W - \widehat{W}^{(N)}) + \frac{1}{2} D^2 F(\widehat{W}^{(N)}) \cdot (W - \widehat{W}^{(N)})^{\otimes 2} \\ &\quad + \frac{1}{6} D^3 F(\zeta) \cdot (W - \widehat{W}^{(N)})^{\otimes 3}, \quad \zeta \in (\widehat{W}^{(N)}, W), \\ \mathbb{E} F(W) &= \mathbb{E} F(\widehat{W}^{(N)}) + \frac{c}{2 \log N} + o \left((\log N)^{-\frac{3}{2} + \varepsilon} \right). \end{aligned}$$

Then, one can design a log-Romberg extrapolation by considering N, N' , $N < N'$ (e.g. $N' \approx 4N$), so that

$$\mathbb{E}(F(W)) = \frac{\log N' \times \mathbb{E}(F(\widehat{W}^{(N')})) - \log N \times \mathbb{E}(F(\widehat{W}^{(N)}))}{\log N' - \log N} + o \left((\log N)^{-\frac{3}{2} + \varepsilon} \right).$$

For practical implementation, it is suggested in [56] to replace $\log N$ by the more consistent “estimator” $\|W - \widehat{W}^{(N)}\|_2^{-2}$.

In fact Assumption (16) holds true for optimal product quantization when F is polynomial function F , $d^0 F = 2$. Assumption (17) holds true in that case as well (see [21]). As concerns optimal quantization, these statements are still conjectures. However, given that \widehat{W} and $W - \widehat{W}$ are independent (see [29]), (16) is equivalent to the simple case where $D^2 F(\widehat{W}^{(N)})$ is constant.

Note that the above extrapolation or some variants can be implemented with other stochastic processes in accordance with the rate of convergence of the quantization error.

9.2 Asian option pricing in a Heston stochastic volatility model

In this section, we will price an Asian call option in a Heston stochastic volatility model using some optimal (at least optimized) functional quantization of the two Brownian motions that drive the diffusion. This model has already been considered in [44] in which functional quantization was implemented for the first time with some product quantizations of the Brownian motions. The Heston stochastic volatility model was introduced in [22] to model stock price

dynamics. Its popularity partly comes from the existence of semi-closed forms for vanilla European options, based on inverse Fourier transform and from its ability to reproduce some skewness shape of the implied volatility surface. We consider it under its risk-neutral probability measure.

$$\begin{aligned} dS_t &= S_t(r dt + \sqrt{v_t} dW_t^1), \quad S_0 = s_0 > 0, \quad (\text{risky asset}) \\ dv_t &= k(a - v_t)dt + \vartheta \sqrt{v_t} dW_t^2, \quad v_0 > 0 \text{ with } d\langle W^1, W^2 \rangle_t = \rho dt, \quad \rho \in [-1, 1]. \end{aligned}$$

where ϑ, k, a such that $\vartheta^2/(4ak) < 1$. We consider the Asian Call payoff with maturity T and strike K . No closed form is available for its premium

$$\text{AsCall}^{Hest} = e^{-rT} \mathbb{E} \left(\frac{1}{T} \int_0^T S_s ds - K \right)^+.$$

We briefly recall how to proceed (see [44] for details): first, one projects W^1 on W^2 so that $W^1 = \rho W^2 + \sqrt{1 - \rho^2} \widetilde{W}^1$ and

$$\begin{aligned} S_t &= s_0 \exp \left(\left(r - \frac{1}{2} \bar{v}_t \right) t + \rho \int_0^t \sqrt{v_s} dW_s^2 \right) \exp \left(\sqrt{1 - \rho^2} \int_0^t \sqrt{v_s} d\widetilde{W}_s^1 \right) \\ &= s_0 \exp \left(t \left(r - \frac{\rho a k}{\vartheta} \right) + \bar{v}_t \left(\frac{\rho k}{\vartheta} - \frac{1}{2} \right) \right) + \frac{\rho}{\vartheta} (v_t - v_0) \exp \left(\sqrt{1 - \rho^2} \int_0^t \sqrt{v_s} d\widetilde{W}_s^1 \right). \end{aligned}$$

The chaining rule for conditional expectations yields

$$\text{AsCall}^{Hest}(s_0, K) = e^{-rT} \mathbb{E} \left(\mathbb{E} \left(\left(\frac{1}{T} \int_0^T S_s ds - K \right)^+ \middle| \sigma(W_t^2, 0 \leq t \leq T) \right) \right).$$

Combining these two expressions and using that \widetilde{W}^1 and W^2 are independent show that $\text{AsCall}^{Hest}(s_0, K)$ is a functional of (\widetilde{W}_t^1, v_t) (as concerns the squared volatility process v , only v_T and $\int_0^T v_s ds$ are involved).

Let $\Gamma^N = \{w_1^N, \dots, w_N^N\}$ be an N -quantizer of the Brownian motion. One solves for $i = 1, \dots, N$, the differential equations for (v_t)

$$dy_i(t) = k \left(a - y_i(t) - \frac{\vartheta^2}{4k} \right) dt + \vartheta \sqrt{y_i(t)} dw_i^N(t), \quad y_i(0) = v_0, \quad (18)$$

using *e.g.* a Runge-Kuta scheme. Let $y_i^{n,N}$ denote the approximation of y_i resulting from the resolution of the above ODE_i ($1/n$ is the time discretization parameter of the scheme). Set the (non-Voronoi) N -quantization of (v_t, S_t) by

$$\widetilde{v}_t^{n,N} = \sum_i y_i^{n,N}(t) \mathbf{1}_{C_i(\Gamma^N)}(W^2) \quad (19)$$

$$\widetilde{S}_t^{n,N} = \sum_{1 \leq i, j \leq N} s_{i,j}^{n,N}(t) \mathbf{1}_{C_i(\Gamma^N)}(\widetilde{W}^1) \mathbf{1}_{C_j(\Gamma^N)}(W^2) \quad (20)$$

$$\begin{aligned} \text{with } s_{i,j}^{n,N}(t) &= s_0 \exp \left(t \left(r - \frac{\rho a k}{\vartheta} \right) + \bar{y}_j^{n,N}(t) \left(\frac{\rho k}{\vartheta} - \frac{1}{2} \right) + \frac{\rho}{\vartheta} (y_j^{n,N}(t) - v_0) \right) \\ &\quad \times \exp \left(\sqrt{1 - \rho^2} \int_0^t \sqrt{y_j^{n,N}(s)} dw_i^N(s) \right) \\ \text{and } \bar{y}_j^{n,N}(t) &= \int_0^t y_j^{n,N}(s) ds. \end{aligned}$$

Note this formula requires the computation of a quantized stochastic integral $\int_0^t \sqrt{y_j^{n,N}(s)} dw_i^N(s)$ (which corresponds to the independent case).

The weights of the product cells $\{\widetilde{W}^1 \in C_i(\Gamma^N), W^2 \in C_j(\Gamma^N)\}$ is given by

$$\mathbb{P}(\widetilde{W}^1 \in C_i(w^N), W^2 \in C_j(w^N)) = \mathbb{P}(\widetilde{W}^1 \in C_i(\Gamma^N)) \mathbb{P}(W^2 \in C_j(\Gamma^N))$$

owing to the independence. For practical implementations different sizes of quantizers can be considered to quantize \widetilde{W}^1 and W^2 .

We follow the guidelines of the methodology introduced in [44]: we compute the *crude* quantized premium for two sizes N and N' , then proceed a space Romberg log-extrapolation. Finally, we make a K -linear interpolation based on the (Asian) forward moneyness $s_0 e^{rT} \frac{1 - e^{-rT}}{rT} \approx s_0 e^{rT}$ (like in [44]) and the Asian Call-Put parity formula

$$\text{AsianCall}^{Hest}(s_0, K) = \text{AsianPut}^{Hest}(s_0, K) + s_0 \frac{1 - e^{-rT}}{rT} - K e^{-rT}.$$

The *anchor strikes* K_{\min} and K_{\max} of the extrapolation are chosen symmetric with respect to the forward moneyness. At K_{\max} , the Call is deep out-of-the-money: one uses the Romberg extrapolated FQ computation; at K_{\min} the Call is deep in-the-money: one computes the Call by parity. In between, one proceeds a linear interpolation in K (which yields the best results, compared to other extrapolations like the quadratic regression approach).

◦ *Parameters of the Heston model:* $s_0 = 100$, $k = 2$, $a = 0.01$, $\rho = 0.5$, $v_0 = 10\%$, $\vartheta = 20\%$.

◦ *Parameters of the option portfolio:* $T = 1$, $K = 99, \dots, 111$ (13 strikes).

◦ *The reference price* has been computed by a 10^8 trial Monte Carlo simulation (including a time Romberg extrapolation of the Euler scheme with $2n = 256$).

◦ *The differential equations (18)* are solved with the parameters of the quantization cubature formulae $\Delta t = 1/32$, with couples of quantization levels $(N, M) = (400, 100), (1000, 100), (3200, 400)$.

Functional Quantization can compute a whole vector (more than 10) option premia for the Asian option in the Heston model with **1 cent accuracy in less than 1 second** (implementation in *C* on a 2.5 GHz processor).

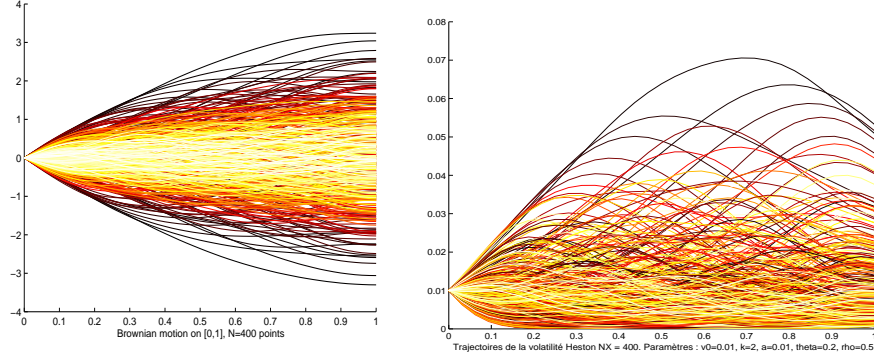


Fig. 10. N -quantizer of the Heston squared volatility process (v_t) ($N = 400$) resulting from an (optimized) N -quantizer of W .




Fig. 11. Quantized diffusions based on optimal functional quantization: Pricing by K -Interpolated-log-Romberg extrapolated- FQ prices as a function of K : absolute error with $(N, M) = (400, 100)$, $(N, M) = (1000, 100)$, $(N, M) = (3200, 400)$. $T = 1$, $s_0 = 50$, $K \in \{99, \dots, 111\}$. $k = 2$, $a = 0.01$, $\rho = 0.5$, $\vartheta = 0.1$.

Further numerical tests carried out or in progress with the B - S model and with the $SABR$ model (Asian, vanilla European options) show the same efficiency. Furthermore, recent attempt to quantize the volatility process and the asset dynamics at different level of quantizations seem very promising in two directions: reduction of the computation time and increase of the robustness of the method to parameter change.

9.3 Comparison: optimized quantization *vs* (optimal) product quantization

The comparison is balanced and probably needs some further *in situ* experiments since it may depend on the modes of the computation. However, it




Fig. 12. Quantized diffusions based on optimal functional quantization: Pricing by K -Interpolated-log-Romberg extrapolated-FQ price as a function of K : convergence as $\Delta t \rightarrow 0$ with $(N, M) = (3200, 400)$ (absolute error). $T = 1$, $s_0 = 50$, $K \in \{99, \dots, 111\}$. $k=2$, $a=0.01$, $\rho=0.5$, $\vartheta=0.1$.




Fig. 13. Quantized diffusions based on optimal product quantization: Pricing by K -linear interpolation of Romberg log-extrapolations as a function of K (absolute error) with $(M, N) = (96, 966)$, $(966, 9984)$. $T = 1$, $s_0 = 50$, $k = 2$, $a = 0.01$, $\rho = 0.5$, $\vartheta = 0.1$. $K \in \{44, \dots, 56\}$.

seems that product quantizers (as those implemented in [44]) are from 2 up to 4 times less efficient than optimal quantizers within our range of application (small values of N). On the other hand, the design of product quantizer from 1-dim scalar quantizers is easy and can be made from some light elementary “bricks” (the scalar quantizer up to $N = 35$ and the optimal allocation rules). Thus, the whole set of data needed to design all optimal product quantizers up to $N = 10\,000$ is approximately 500 KB whereas one optimal quantizer with size $10\,000 \approx 1$ MB...

10 Universal quantization rate and mean regularity

The following theorem points out the connection between functional quantization rate and mean regularity of $t \mapsto X_t$ from $[0, T]$ to $L^r(\mathbb{P})$.

Theorem 8. ([33] (2005)) *Let $X = (X_t)_{t \in [0, T]}$ be a stochastic process. If there is $r^* \in (0, \infty)$ and $a \in (0, 1]$ such that*

$$X_0 \in L^{r^*}(\mathbb{P}), \quad \|X_t - X_s\|_{L^{r^*}(\mathbb{P})} \leq C_X |t - s|^a,$$

for some positive real constant $C_X > 0$, then

$$\forall p, r \in (0, r^*), \quad e_{N,r}(X, L_T^p) = O((\log N)^{-a}).$$

The proof is based on a constructive approach which involves the Haar basis (instead of K - L basis), the non asymptotic version Zador Theorem and product functional quantization. Roughly speaking, we use the unconditionality of the Haar basis in every L_T^p (when $1 < p < \infty$) and its wavelet feature *i.e.* its ability to “code” the path regularity of a function on the decay rate of its coordinates.

EXAMPLES (SEE [33]): • d -dimensional Itô processes (includes d -dim diffusions with sublinear coefficients) with $a = 1/2$.

• General Lévy process X with Lévy measure ν with square integrable big jumps. If X has a Brownian component, then $a = 2$, otherwise if $\beta(X) > 0$ where $\beta(X) := \inf \{ \theta : \int |y|^\theta \nu(dy) < +\infty \} \in (0, 2)$ (Blumenthal-Gettoor index of X), then $a = \beta^*(X)$. This rate is the *exact rate i.e.*

$$e_{N,r}(X, L_T^p) \approx (\log N)^{-a}$$

for many classes of Lévy processes like symmetric stable processes, Lévy processes having a Brownian component, etc (see [33] for further examples).

• When X is a compound Poisson processes, then $\beta(X) = 0$ and one shows, still with constructive methods, that

$$e_N(X) = O(e^{-(\log N)^\vartheta}), \quad \vartheta \in (0, 1),$$

which is in-between the finite and infinite dimensional settings.

11 About lower bounds

In this overview, we gave no clue toward lower bounds although most of the rates we mentioned are either exact (\approx) or sharp (\sim) (we tried to emphasize the numerical aspects). Several approaches can be developed to get some lower bounds. Historically, the first one was to rely on subadditivity property of the quantization error derived from self-similarity of the distribution: this works

with the uniform distribution over $[0, 1]^d$ but also in an infinite dimensional framework (see *e.g.* [12] for the fractional Brownian motion).

A second approach consists in pointing out the connection with the Shannon-Kolmogorov entropy (see *e.g.* [29]) using that the entropy of a random variable taking at most N values is at most $\log N$.

A third connection can be made with small deviation theory (see [9], [19] and [33]). Thus, in [19], a connection is established between (functional) quantization and small ball deviation for Gaussian processes. In particular this approach provides a method to derive a lower bound for the quantization rate from some upper bound for the small deviation problem. A careful reading of the proof of Theorem 1.2 in [19] shows that this small deviation lower bound holds for any *unimodal* (w.r.t. 0) non zero process. To be precise: assume that \mathbb{P}_X is L_T^p -unimodal *i.e.* there exists a real $\varepsilon_0 > 0$ such that

$$\forall x \in L_T^p, \forall \varepsilon \in (0, \varepsilon_0], \quad \mathbb{P}(|X - x|_{L_T^p} \leq \varepsilon) \leq \mathbb{P}(|X|_{L_T^p} \leq \varepsilon).$$

For centered Gaussian processes (or processes “subordinated” to Gaussian processes) this follows from the Anderson Inequality (when $p \geq 1$). If

$$G(-\log(\mathbb{P}(|X|_{L_T^p} \leq \varepsilon))) = \Omega(1/\varepsilon) \quad \text{as } \varepsilon \rightarrow 0$$

for some increasing unbounded function $G : (0, \infty) \rightarrow (0, \infty)$, then

$$\forall c > 1, \quad \liminf_N G(\log(cN))e_{N,r}(X, L_T^p) > 0, \quad r \in (0, \infty). \quad (21)$$

This approach is efficient in the non quadratic case as emphasized in [33] where several universal bounds are shown to be optimal using this approach.

ACKNOWLEDGEMENT. I thank S. Graf, H. Luschgy J. Printems and B. Wilbertz for all the fruitful discussions and collaborations we have about functional quantization.

References

1. ABAYA, E.F. AND WISE, G.L. (1982). On the existence of optimal quantizers. *IEEE Trans. Inform. Theory*, **28**, 937-940
2. ABAYA, E.F. AND WISE, G.L. (1984). Some remarks on the existence of optimal quantizers. *Statistics and Probab. Letters*, **2**, 349-351.
3. BENVENISTE, A., MÉTIVIER, M. AND PRIOURET, P. (1990). *Adaptive algorithms and stochastic approximations*, Translated from the French by Stephen S. Wilson. Applications of Mathematics **22**, Springer-Verlag, Berlin, 365 pp.
4. N. BOULEAU, D. LÉPINGLE (1994). *Numerical methods for stochastic processes*, Wiley Series in Probability and Mathematical Statistics: Applied Probability and Statistics. A Wiley-Interscience Publication. John Wiley & Sons, Inc., New York, 359 pp. ISBN: 0-471-54641-0.

5. BUCKLEW, J.A. AND WISE, G.L. (1982). Multidimensional asymptotic quantization theory with r^{th} power distortion. *IEEE Trans. Inform. Theory*, **28**(2), 239-247.
6. COHORT, P. (1998). A geometric method for uniqueness of locally optimal quantizer. Pre-print LPMA-464 and Ph.D. Thesis, *Sur quelques problèmes de quantification*, 2000, Univ. Paris 6.
7. CUESTA-ALBERTOS, J.A., MATRÁN, C. (1988). The strong law of large numbers for k -means and best possible nets of Banach valued random variables, *Probab. Theory Rel. Fields* **78**, 523-534.
8. DELATTRE, S., FORT, J.-C. AND PAGÈS, G. (2004). Local distortion and μ -mass of the cells of one dimensional asymptotically optimal quantizers, *Communications in Statistics*, **33**(5), 1087-1118.
9. DEREICH, S., FEHRINGER, F., MATOUSSI, A. AND SCHEUTZOW, M. (2003). On the link between small ball probabilities and the quantization problem for Gaussian measures on Banach spaces, *J. Theoretical Probab.*, **16**, pp.249-265.
10. DEREICH, S. (2005). The coding complexity of diffusion processes under $L^p[0, 1]$ -norm distortion, pre-print.
11. DEREICH, S. (2005). The coding complexity of diffusion processes under supremum norm distortion, pre-print.
12. DEREICH, S., SCHEUTZOW, M. (2006). High resolution quantization and entropy coding for fractional Brownian motions, *Electron. J. Probab.*, **11**, 700-722.
13. DOSS H. (1977). Liens entre équations différentielles stochastiques et ordinaires, *Ann. I.H.P.*, section B, **13**(2), 99-125.
14. FLEISCHER, P.E. (1964). Sufficient conditions for achieving minimum distortion in a quantizer. *IEEE Int. Conv. Rec.*, part I, 104-111.
15. GERSHO, A. AND GRAY, R.M. (1992). *Vector Quantization and Signal Compression*. Kluwer, Boston.
16. FORT, J.-C. AND PAGÈS, G. (2004). Asymptotics of optimal quantizers for some scalar distributions, *Journal of Computational and Applied Mathematics*, **146**, 253-275, 2002.
17. GRAF, S. AND LUSCHGY, H. (2000). *Foundations of Quantization for Probability Distributions*. Lect. Notes in Math. 1730, Springer, Berlin, 230p.
18. GRAF, S. AND LUSCHGY, H. (2005). The point density measure in the quantization of self-similar probabilities. *Math. Proc. Cambridge Phil. Soc.* **138**, 513-531.
19. GRAF, S., LUSCHGY H. AND PAGÈS, G. (2003). Functional quantization and small ball probabilities for Gaussian processes, *J. Theoret. Probab.*, **16**(4), 1047-1062.
20. GRAF, S., LUSCHGY, H., PAGÈS, G. (2007). Optimal quantizers for Radon random vectors in a Banach space, *J. of Approximation Theory*, **144**, 27-53.
21. GRAF, S., LUSCHGY, H. AND PAGÈS, G. (2006). Distortion mismatch in the quantization of probability measures, to appear in *ESAIM P&S*.
22. HESTON, S.L. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options, *The review of Financial Studies*, **6**(2), 327-343.
23. KIEFFER, J.C. (1983). Uniqueness of locally optimal quantizer for log-concave density and convex error weighting functions, *IEEE Trans. Inform. Theory*, **29**, 42-47.
24. KIEFFER, J.C. (1982). Exponential rate of convergence for Lloyd's Method I, *IEEE Trans. Inform. Theory*, **28**(2), 205-210.

25. KUSHNER, H. J., YIN, G. G. (2003). *Stochastic approximation and recursive algorithms and applications*. Second edition. Applications of Mathematics **35**. Stochastic Modelling and Applied Probability. Springer-Verlag, New York, 474p.
26. LAMBERTON, D. AND PAGÈS, G. (1996). On the critical points of the 1-dimensional Competitive Learning Vector Quantization Algorithm. *Proceedings of the ESANN'96*, (ed. M. Verleysen), Editions D Facto, Bruxelles, 97-106.
27. LAPEYRE, B., SAB, K. AND PAGÈS, G. (1990). Sequences with low discrepancy. Generalization and application to Robbins-Monro algorithm, *Statistics*, **21**(2), 251-272.
28. LEJAY, A. (2003). An introduction to rough paths, *Séminaire de Probabilités XXXVII*, Lecture Notes in Mathematics 1832, Springer, Berlin, 1-59.
29. LUSCHGY, H., PAGÈS, G. (2002). Functional quantization of Gaussian processes, *Journal of Functional Analysis*, **196**(2), 486-531.
30. LUSCHGY, H., PAGÈS, G. (2004). Sharp asymptotics of the functional quantization problem for Gaussian processes, *The Annals of Probability*, **32**(2), 1574-1599.
31. LUSCHGY, H., PAGÈS, G. (2006). Functional quantization of a class of Brownian diffusions: A constructive approach, *Stochastic Processes and Applications*, **116**, 310-336.
32. LUSCHGY, H., PAGÈS, G. (2005). High-resolution product quantization for Gaussian processes under sup-norm distortion, pre-pub LPMA-1029, forthcoming in *Bernoulli*.
33. LUSCHGY, H., PAGÈS, G. (2006). Functional Quantization Rate and mean regularity of processes with an application to Lévy Processes, pre-print LPMA-1048.
34. LUSCHGY, H., PAGÈS, G. (2007). Expansion of Gaussian processes and Hilbert frames, technical report.
35. LUSCHGY, H., PAGÈS, G. AND WILBERTZ, B. (2007). Asymptotically optimal quantization schemes for Gaussian processes, in progress.
36. MRAD, M., BEN HAMIDA, S. (2006). Optimal Quantization: Evolutionary Algorithm vs Stochastic Gradient, *Proceedings of the 9th Joint Conference on Information Sciences*.
37. NEWMAN, D.J. (1982). The Hexagon Theorem. *IEEE Trans. Inform. Theory*, **28**, 137-138.
38. H. NIEDERREITER (1992) *Random Number Generation and Quasi-Monte Carlo Methods*, CBMS-NSF regional conference series in Applied mathematics, SIAM, Philadelphia.
39. PAGÈS, G. (1993). Voronoi tessellation, space quantization algorithm and numerical integration. *Proceedings of the ESANN'93*, M. Verleysen Ed., Editions D Facto, Bruxelles, 221-228.
40. PAGÈS, G. (1997). A space vector quantization method for numerical integration, *J. Computational and Applied Mathematics*, **89**, 1-38.
41. PAGÈS, G. (2000). Functional quantization: a first approach, pre-print CMP12-04-00, Univ. Paris 12.
42. PAGÈS, G., PHAM, H. AND PRINTEMS, J. (2003). Optimal quantization methods and applications to numerical methods in finance. *Handbook of Computational and Numerical Methods in Finance*, S.T. Rachev ed., Birkhäuser, Boston, 429p.
43. PAGÈS, G., PRINTEMS, J. (2003). Optimal quadratic quantization for numerics: the Gaussian case, *Monte Carlo Methods and Appl.*, **9**(2), 135-165.
44. PAGÈS, G., PRINTEMS, J. (2005). Functional quantization for numerics with an application to option pricing, *Monte Carlo Methods and Appl.*, **11**(4), 407-446.

45. PAGÈS, G., PRINTEMS, J. (2005). Website devoted to vector and functional optimal quantization: www.quantize.maths-fi.com.
46. PAGÈS, G., SELLAMI, A. (2007). Convergence of multi-dimensional quantized SDE's. In progress.
47. G. PAGÈS, Y.J. XIAO (1988) Sequences with low discrepancy and pseudo-random numbers: theoretical results and numerical tests, *J. of Statist. Comput. Simul.*, **56**, 163-188.
48. PÄRNA, K. (1990). On the existence and weak convergence of k -centers in Banach spaces, *Tartu Ülikooli Toimetised*, **893**, 17-287.
49. POLLARD, D. (1982). Quantization and the method of k -means. *IEEE Trans. Inform. Theory*, **28**(2), 199-205.
50. P.D. PROINOV (1988). Discrepancy and integration of continuous functions, *J. of Approx. Theory*, **52**, 121-131.
51. REVUZ, D., YOR, M. (1999). *Continuous martingales and Brownian motion*, Third edition. Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences], 293, Springer-Verlag, Berlin, 1999, 602 p.
52. K.F. ROTH (1954). On irregularities of distributions, *Mathematika*, **1**, 73-79.
53. TARPEY, T., KINATEDER, K.K.J. (2003). Clustering functional data, *J. Classification*, **20**, 93-114.
54. TARPEY, T., PETKOVA, E., OGDEN, R.T. (2003). Profiling Placebo responders by self-consistent partitioning of functional data, *J. Amer. Statist. Association*, **98**, 850-858.
55. TRUSHKIN, A.V. (1982). Sufficient conditions for uniqueness of a locally optimal quantizer for a class of convex error weighting functions, *IEEE Trans. Inform. Theory*, **28**(2), 187-198.
56. WILBERTZ, B. (2005). Computational aspects of functional quantization for Gaussian measures and applications, diploma thesis, Univ. Trier.
57. ZADOR, P.L. (1963). Development and evaluation of procedures for quantizing multivariate distributions. Ph.D. dissertation, Stanford Univ.
58. ZADOR, P.L. (1982). Asymptotic quantization error of continuous signals and the quantization dimension. *IEEE Trans. Inform. Theory*, **28**(2), 139-149.